

Analog Dialogue


A forum for the exchange of circuits, systems, and software for real-world signal processing

OTHELLO™: A NEW DIRECT-CONVERSION RADIO CHIPSET ELIMINATES IF STAGES (page 3)

Logarithmic Amplifiers Explained—Ask the Applications Engineer—28 (page 30)

All-Electronic Power and Energy Meters (page 64)

Complete contents on page 1



About *Analog Dialogue*

Analog Dialogue is the free technical magazine of Analog Devices, Inc., published continuously for thirty-three years, starting in 1967. It discusses products, applications, technology and techniques for analog, digital, and mixed-signal processing.

Volume 33, the current issue, incorporates all articles published during 1999 in the World Wide Web¹ editions. All recent issues, starting with Volume 29, Number 2 (1995) have been archived on that website.

As one of the many informative Analog Devices publications, *Analog Dialogue*'s objectives are to inform engineers, scientists, and electronic technicians about new ADI products and technologies, and to help them understand and competently apply our products.

The monthly Web editions have at least three further objectives:

- Provide timely digests that alert readers to upcoming and available products.
- Provide a set of links to important and rapidly proliferating sources of information and activity fermenting within the ADI Web site²
- Listen to reader suggestions and help find sources of aid to answer their questions.

Thus, *Analog Dialogue* is more than a magazine: its links and tendrils to all parts of our website (and some outside sites) make its bookmark a favorite "high-pass-filtered" point of entry to the **analog.com** site—the world of Analog Devices.

Our hope is that readers will think of ADI publications as "Great Stuff" and the *Analog Dialogue* bookmark on their web browser as a favorite alternative path to answer the question, "What's new at ADI?"

Welcome! Read and enjoy!



Dan Sheingold
dan.sheingold@analog.com
Editor, *Analog Dialogue*

¹www.analog.com/analogdialogue/

²www.analog.com

IN THIS ISSUE

Analog Dialogue Volume 33, 1999

	Page
Editor's Notes, Authors	2
<i>Othello: A new direct-conversion radio chip set eliminates IF stages</i>	3
<i>DSPs enhance flexible third-generation (3G) base station design</i>	6
<i>Phase-locked loops for high-frequency receivers and transmitters, Part 1</i>	9
<i>Phase-locked loops for high-frequency receivers and transmitters, Part 2</i>	13
<i>Phase-locked loops for high-frequency receivers and transmitters, Part 3</i>	18
<i>Dual-axis, low-g, fully integrated accelerometers</i>	23
<i>Accelerometers—fantasy and reality (Ask The Applications Engineer—29)</i>	25
<i>ADXL105: A lower-noise, wider-bandwidth accelerometer, rivals performance of more expensive sensors</i>	27
<i>Logarithmic amplifiers explained (Ask The Applications Engineer—28)</i>	30
<i>Analog-to-Digital converter architectures and choices for system design</i>	35
New-Product Brief: AD9814 Low Power 14-Bit, 3-Channel CCD Signal Processor	38
New-Product Brief: AD9884 8-bit, 140 MSPS Flat Panel Display Interface	38
<i>New TxDAC® Generation: 125 MSPS 10-, 12-, and 14-bit high-performance DACs for wideband multitone communication Transmit channels</i>	39
New-Product Brief: AD9772 14-bit, 300 MSPS TxDAC	40
<i>Ultrasound analog electronics primer</i>	41
New-Product Brief: ADP3421 Geyserville-Enabled DC/DC Converter Controller	43
<i>A chip you can use to monitor environmental conditions on PC motherboard designs</i>	44
<i>Measuring Temperatures on Computer Chips with Speed and Accuracy: A new approach using silicon sensors and off-chip processing</i>	49
New-Product Brief: AD1881 AC'97 SoundMAX® Audio Codec	53
<i>Compensate for loading effects on power lines with a DSP-controlled active shunt filter</i>	54
<i>Signal corruption in industrial measurement (Ask The Applications Engineer—27)</i>	58
New-Product Brief: AD8551/2/4 Zero-Drift, Single-Supply, Rail-to-Rail Op Amps	59
<i>Process signals from millivolts to ± 10 V directly with a versatile single-supply 3/5V charge-balancing A/D converter</i>	60
New-Product Brief: ADuC812 12-bit MicroConverter Data Acquisition System	61
<i>Microcontroller-based energy metering using the AD7755</i>	62
New-Product Brief: ADMC401 DSP-Based High-Performance Motor Controller	63
<i>All-electronic power and energy meters</i>	64
New-Product Brief: AD7751 Fault-Tolerant Energy Metering IC	66
<i>Make vs. Buy: When should I re-invent the wheel?</i>	67
New ADI Fellows	70
Authors (continued from page 2)	71
Worth Reading: An authoritative, practical, readable new DSP book. Read it FREE on-line.	72

Cover: The cover illustration was designed and executed by Kristine Chmiel-Lafleur, of Communications Services, Analog Devices, Inc.

Editor's Notes

1999: THE YEAR THAT WAS

In 1999, with the aid of the World Wide Web, *Analog Dialogue* has increased its value to its readers and simultaneously increased its value to its sponsor, Analog Devices, Inc. With a target of monthly publication, ten issues were published in 1999, more than three times as many as the yearly average over its prior 32 years in print. The greater frequency of publication made it possible for readers to read about new products in timely fashion, and the ease of publishing on the Web allowed us to literally "turn on a dime". More importantly, though, it allowed readers immediately to link to extensive further information on products and topics of interest on the vast Analog Devices Web Site and to promptly print usable data sheets. This essentially eliminated the need for Business Reply cards and the lengthy "circle and wait" cycle.

The one downside and source of anguish to many faithful readers was that, in order to read *Analog Dialogue*, they had to access the Web (a task that day by day is nevertheless becoming easier with the help of some ADI broadband and DSL products). This issue, Volume 33, divorced from the quest for immediacy that energizes our Web version, is our answer to the need for a "permanent" version of *Analog Dialogue* that you can hold in your hand. It contains all the technical articles (and a few worthy New-Product Briefs) that have been published during 1999. You can carry it, read it wherever you are (except in the dark), and stack it in a binder with previous issues to complete your set. And, in between this edition and Volume 34, in December, 2000, you can still get up-to-date information from the monthly editions of ADI's technical magazine on the Web.¹ Surely this is the best of both worlds!

SCOTT WAYNE

Contributing Editor

We're pleased to welcome Scott Wayne to the staff of *Analog Dialogue* in the role of Contributing Editor. Scott is the newest addition to the Corporate Technical Communications group at Analog Devices.


Besides contributing counsel and articles to this publication, Scott works with engineers at all levels and locations within Analog Devices, with editors of trade publications, and with a wide range of technical people throughout the world to develop interesting and useful stories about designs, technologies, and applications of Analog Devices products. He also has a role in one of the more important aspects of professional development—encouraging, stimulating, and helping engineers to overcome their reticence to write technical articles.

Scott has an SBEE from MIT. He interned with Analog Devices in 1978, and came back after graduating in 1979. Since then, he has designed high-resolution analog-to-digital converters, digital-to-analog converters, sample-and-hold amplifiers, instrumentation amplifiers, isolation amplifiers, and an iontophoretic drug-delivery system using modular, hybrid, compound-monolithic, and



monolithic technologies. In June of 1999, Scott transferred out of Design, and into Corporate Marketing, where he is ADI's Senior Technical Editor.

Scott has written articles and design ideas that have appeared in the trade press, and he was a contributor to the 1986 edition of the *Analog-Digital Conversion Handbook*. He holds a patent for a precision switched-capacitor-ratio system and has a second patent pending. He has presented several times at the Analog Devices General Technical Conference.

In his free time, Scott enjoys sailing, hiking, golf, bicycling, watching Boston's Red Sox and Bruins, and hanging out with his "Little Brother," Gary. Dan.Sheingold@analog.com 

THE AUTHORS


Brian Black (page 35), who joined Analog Devices in 1998, is Product Manager, high-resolution data converters, for the General-Purpose Converter group, in Wilmington, MA. He has a BSEE from Brigham Young University (1991), and MS degrees from MIT (1996) in both EE and Management (Sloan School).




Michael Curtin (pages 9-22) is a Staff Applications Engineer at our Limerick, Ireland facility, providing application support for the frequency-synthesis products. Mike was graduated from the University of Limerick with a BSc. He is the author of several technical articles on high-resolution A/D and D/A converters and is the holder of two patents. For relaxation, he enjoys playing snooker, reading, and walking the family dog.



Robert De Robertis (page 6) manages the Wireless Infrastructure Business Group in ADI's DSP division, which includes applications of products such as TigerSHARC and the ADSP-219x. His activities in the past decade have centered on leveraging DSP technologies into emerging digital wireless markets. In his various engineering, marketing, and product-line management positions at Analog Devices, Cadence Design Systems, Lucent Technologies, and Harris Semiconductor, Rob has experienced many aspects of the wireless revolution.



Paul Daigle (page 64) joined Analog Devices in 1997 as a Product Manager in the Standard Linear Products division. He holds an MSEE and BSEE from Syracuse University. Paul has been participating in the semiconductor industry since 1989. He thoroughly enjoys life as father to a 1-year-old son and a 3-year-old daughter. Mountain hikes and seaside walks are his family's favorites.



[more authors on Page 71]

¹<http://www.analog.com/analogdialogue/>

Othello: A New Direct-Conversion Radio Chip Set Eliminates IF Stages

by Dan Fague

INTRODUCTION

Analog Devices recently announced the revolutionary Othello direct-conversion radio for mobile applications. By eliminating intermediate-frequency (IF) stages, this chip set will permit the mobile electronics industry to reduce the size and cost of radio sections and enable flexible, multistandard, multimode operation. The radio consists of two integrated circuits, the AD6523 Zero-IF Transceiver and the AD6524 Multiband Synthesizer. The AD6523 contains the main functions necessary for both a direct-conversion receiver and a direct VCO transmitter, known as the Virtual-IF™ transmitter. It also includes the local-oscillator generation block and a complete on-chip regulator that supplies power to all active circuitry for the radio. The AD6524 is a fractional-*N* synthesizer that features extremely fast lock times to enable advanced data services over cellular telephones—such as high-speed circuit-switched data (HSCSD) and general packet radio services (GPRS).

Together, the two ICs supply the main functions necessary for implementing dual- or triple-band radios for GSM cellular phones. The direct conversion technology, combined with a new twist on the translation loop (or direct VCO) modulator, reduces the amount of external filtering needed in the radio to an absolute minimum.

THE GSM STANDARD

The Global System for Mobiles (GSM) was officially launched in 1992, after over five years of standards writing by the European Telecommunications Standards Institute (ETSI). The goal of GSM was to unite a Babel of European communications under one digital cellular standard. Before GSM, Europe maintained in effect one separate cellular network for each country, making international roaming on the continent virtually impossible. With GSM, a citizen of any of the original seventeen countries could roam to any other country using a single cellular handset. The standard, which was written with future expansion to data services and other applications in mind, soon became popular around the world. It is now accepted in more than 140 countries, with over 200 networks running.

The frequency bands originally allocated to GSM were 890 to 915 MHz for mobile transmitting and 935 to 960 MHz for mobile receiving. That band was expanded to the so-called E-GSM bands of 880 to 915 MHz and 925 to 960 MHz. Another frequency allocation was made to further expand GSM capacity. This band, allocated to digital communications services (DCS), was 1710 to 1785 MHz and 1805 to 1880 MHz. All countries adopting GSM use one of these two pairs of frequency bands, except the United States, where both bands were already allocated by the FCC. The Personal Communications Services (PCS) frequency auctions in the mid-1990s made available a set of bands for GSM in the U.S.—1850 to 1910 MHz and 1930 to 1990 MHz.

Othello, Superhomodyne and Virtual-IF are trademarks of Analog Devices, Inc.

Today's typical GSM handset (or handy) will have 2-W output power and is required to receive signals as low as -102 dBm (less than 1/10 of a picowatt). The handy includes a powerful digital signal processor (DSP) core (equivalent to an ADSP-218x) to encode, encrypt, interleave, packetize, transmit, receive, de-packetize, de-interleave, de-encrypt, and de-encode the data going to and coming from the voiceband A/D and D/A converters. An equally powerful microcontroller (ARM or Hitachi H8), combined with a hardware burst processor, controls the timing necessary to implement the time-division multiple-access (TDMA) and frequency hopping functions to keep the phone call on a specific time and frequency channel. The microcontroller also implements the man-machine interface, and operates all the necessary protocols for communication to the base stations.

RADIO ARCHITECTURE DESIGN

Most digital cellular phones today include at least one "downconversion" in their signal chain. This frequency conversion shifts the desired signal from the allocated RF band for the standard (say, at 900 MHz) to some lower intermediate frequency (IF), where channel selection is performed with a narrow channel-select filter (usually a surface acoustic-wave (SAW) or a ceramic type). The now-filtered signal is then further down-converted to either a second IF or directly to baseband, where it is digitized and demodulated in a digital signal processor (DSP).

The idea of using direct-conversion for receivers has long been of interest in RF design. The reason is obvious: in consumer equipment conversion stages add cost, bulk, and weight. Each conversion stage requires a local oscillator, (often including a frequency synthesizer to lock the LO onto a given frequency), a mixer, a filter, and (possibly) an amplifier. No wonder, then, that direct conversion receivers would be attractive. All intermediate stages are eliminated, reducing the cost, volume, and weight of the receiver.

The first Othello radio reduces the component count even more by integrating the front-end GSM low-noise amplifier (LNA). This eliminates an RF filter (the "image" filter) that is necessary to eliminate the image, or unwanted mixing product of a mixer and the off chip LNA. This stage, normally implemented with a discrete transistor, plus biasing and matching networks, accounts for a total of about 12 components. Integrating the LNA saves a total of about 15 to 17 components, depending on the amount of matching called for by the (now-eliminated) filter.

SUPERHOMODYNE™ DIRECT-CONVERSION RECEIVER

A functional block diagram of the Othello dual band GSM radio's architecture is shown in Figure 1. The receive section is at the top of the figure. From the antenna connector, the desired signal enters the transmit/receive switch and exits on the appropriate path, either 925-960 MHz for the GSM band or 1805-1880 MHz for DCS. The signal then passes through an RF band filter (a so-called "roofing filter") that serves to pass the entire desired frequency band while attenuating all other out-of-band frequencies (blockers—including frequencies in the transmission band) to prevent them from saturating the active components in the radio front end. The roofing filter is followed by the low-noise amplifier (LNA). This is the first gain element in the system, effectively reducing the contribution of all following stages to system noise. After the LNA, the direct-conversion mixer translates the desired signal from radio frequency (RF) all the way to baseband by

multiplying the desired signal with a local oscillator (LO) output at the same frequency.

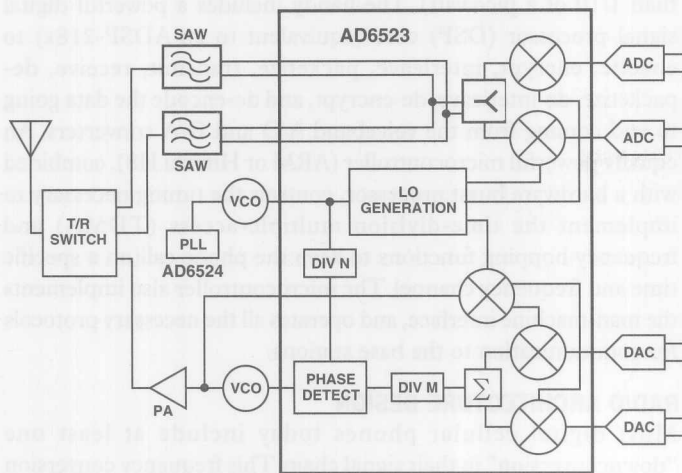


Figure 1. Block diagram of the Othello dual-band radio.

The output of the mixer stage is then sent in quadrature (I and Q channels) to the variable-gain baseband amplifier stage. The VGA also provides some filtering of adjacent channels, and attenuation of in-band blockers. These blocking signals are other GSM channels that are some distance from the desired channel, say 3 MHz and beyond. The baseband amplifiers filter these signals so that they will not saturate the Receive ADCs. After the amplifier stage, the desired signal is digitized by the Receive ADCs.

VIRTUAL-IF™ TRANSMITTER

The Transmit section begins on the right, at the multiplexed I and Q inputs/outputs. Because the GSM system is a time division duplex (TDD) system, the transmitter and receiver are never on at the same time. The Othello radio architecture takes advantage of this fact to save four pins on the transceiver IC's package. The quadrature transmit signals enter the transmitter through the multiplexed I/Os. These I and Q signals are then modulated onto a carrier at an intermediate frequency greater than 100 MHz.

The output of the modulator goes to a phase-frequency detector (PFD), where it is compared to a reference frequency that is generated from the external channel selecting LO. The output of the PFD is a charge pump, operating at above 100 MHz, whose output is filtered by a fairly wide (1 MHz) loop filter. The output of the loop filter drives the tuning port of a voltage-controlled oscillator (VCO), with frequency ranges that cover the GSM and DCS transmit bands.

The output of the transmit VCO is sent to two places. The main path is to the transmit power amplifier (PA), which amplifies the transmit signal from about +3 dBm to +35 dBm, sending it to the transmit/receive switch and low-pass filter (which attenuates power-amplifier harmonics). The power amplifiers are dual band, with a simple CMOS control voltage for the band switch. The VCO output also goes to the transmit feedback mixer by means of a coupler, which is either a printed circuit, built with discrete inductors and capacitors, or a monolithic (normally ceramic) coupling device. The feedback mixer downconverts the transmit signal to the transmit IF, and uses it as the local oscillator signal for the transmit modulator.

This type of modulator has several names, but the most descriptive is probably "translation loop." The translation loop modulator takes

advantage of one key aspect of the GSM standard: the modulation scheme is Gaussian-filtered minimum-shift keying (GMSK). This type of modulation does not affect the envelope amplitude, which means that a power amplifier can be saturated and still not distort the GMSK signal sent through it.

GMSK can be generated in several different ways. In another European standard (for cordless telephones), GMSK is created by directly modulating a free running-VCO with the Gaussian filtered data stream. In GSM, the method of choice has been quadrature modulation. Quadrature modulation creates accurate phase GMSK, but imperfections in the modulator circuit (or up-conversion stages) can produce envelope fluctuations, which can in turn degrade the phase trajectory when amplified by a saturated power amplifier. To avoid such degradations, GSM phone makers have been forced to use amplifiers with somewhat higher linearity, at the cost of reduced efficiency and talk time per battery charge cycle.

The translation loop modulator combines the advantages of directly modulating the VCO and the inherently more accurate quadrature modulation. In effect, the scheme creates a phase locked loop (PLL), comprising the modulator, the LO signal, and the VCO output and feedback mixer. The result is a directly modulated VCO output with a perfectly constant envelope and almost perfect phase trajectory. Phase trajectory errors as low as 1.5 degrees have been measured in the AD6523 transceiver IC, using a signal generator as the LO signal to provide a reference for the loop.

FREQUENCY PLANNING

An important aspect of the Othello radio design is the frequency plan. The GSM standard has strict requirements with regard to in- and out-of-band spurious emissions. A GSM cellular phone must be able to withstand blockers at extremely high levels (0 dBm) while continuing to receive normally. The phone must also not *emit* spurious signals into other bands above a certain level (in the GSM receive band, -112 dB relative to the transmitted signal!).

The Othello radio architecture was designed with the entire system in mind. The frequency plan was carefully crafted to satisfy three equally important criteria:

- 1) Reduce spurious emissions from the radio.
- 2) Minimize bandwidth of the dual band local oscillator (LO) VCO.
- 3) Eliminate as many potential blockers as possible.

By satisfying all of these criteria, major radio problems have been solved, always keeping the end-user and the application in mind. The final solution turned out to be both elegant and practical.

Reducing Spurious Emissions from the Radio

Spurious emissions from the radio can cause problems in both Transmit and Receive modes. A wayward LO signal can find its way to the antenna and "self block" a direct-conversion receiver, reducing sensitivity. The LO signal can also radiate from the antenna and degrade the performance of other receivers.

In the Othello frequency plan, the local oscillator's center frequency was chosen to be about 1350 MHz. This placed the LO strategically between the GSM and DCS frequency bands, enabling a single LO to be used for both GSM and DCS, saving components. Since that frequency is distant from either of the bands, the radio's front end filters will attenuate any radiated LO signal, and so it doesn't pose a problem as a radiated spurious emission. Even if the signal is coupled directly from pin to pin on the IC, its power level would

be lower than the GSM requirements for in- or out-of-band blockers received at the antenna.

In the Transmit section, spurious signals can also pose a problem. Though the transmitter is a direct VCO modulator, the feedback mixer will introduce spurious signals at its output that must be filtered before entering the phase detector. Otherwise, they could appear at the output themselves or cause still other spurious signals to appear by mixing with the desired modulation signal due to the non-linear operation of the phase detector input stage. This is a problem inherent in any translation-loop modulator. By using a widely separated LO frequency, the Othello architecture simplifies filtering of these products.

Minimize Bandwidth of the Dual Band Local-Oscillator VCO

The Othello architecture was designed to minimize the number of external components needed to build a complete dual-band radio. The frequency plan was specifically chosen to make it possible for a single LO VCO to cover both GSM and DCS frequency bands while still meeting the necessarily stringent phase-noise specifications at the 3-MHz offset demanded of all GSM LO VCOs. By keeping the bandwidth requirements of the VCO to a minimum, the VCO can be designed with a maximum supply voltage of 2.7 V. This allows the entire dual-band radio to run at 2.7 V, reducing power consumption and enabling the use of nickel-cadmium (NiCd), nickel metal-hydride (NiMH), or lithium-ion (Li-ion) battery types.

Eliminate as Many Potential Blockers as Possible

As a result of the direct-conversion receiver architecture, the Othello radio has fewer "trouble" channels for the blocking tests required by GSM. Superheterodyne receivers must always contend with half-IF responses that are difficult to filter with RF filters due to the shape factors required. By going to direct conversion Othello eliminates the half-IF response.

PERFORMANCE

One of the key advantages of the Othello radio is that the reduction in the number of components needed to implement it does not engender performance sacrifices. In both the GSM and DCS bands, the Othello system noise figure allows for a production margin of about 6 dB from the required receiver sensitivity of -102 dBm. The transmitter provides a similar production margin, with phase trajectory errors of 2.5° rms, compared with the requirement of 5° rms.

FUTURE BENEFITS

Another important feature of the Othello radio is that the AD6524's fractional-N synthesizer has a lock time short enough to enable GPRS operation. [GPRS, an extension to the GSM network coming in year 2000, will allow very high data rates to be used by a compliant GSM handset.] A requirement of GPRS operation is that the LO synthesizer must lock in less than half a GSM time slot (lock times less than 250 μ s). The AD6524, with its fractional-N synthesizer, is able to reach lock faster than conventional synthesizers, because fractional-N types operate at reference frequencies that are higher than the channel spacing, thus jumping more than one channel per reference cycle. In the case of the AD6524, the 26-MHz reference frequency, twice the system crystal frequency (compared to a channel spacing of 200 kHz), ensures that the Othello radio will meet the required lock time for GPRS. Fast lock time also helps to reduce power consumption by allowing the baseband section to keep the radio off for longer time intervals.

The Othello radio has opened new doors of opportunity for the future. Today, a complete dual-band Othello radio, including all power-management functions, can be implemented with only 90 components. With so few components, the radio can be implemented in less than 10 cm² of board space. Figure 2 is a photograph of a prototype radio design for Othello, implemented on a four layer PCB. Compare this to a superheterodyne receiver that today uses about 225 components crowded into somewhat less than 15 cm² of board space for the same functionality. (Even that is an improvement over radios of just two years ago, which used the same number of components to implement a single band GSM radio!) The advantages of direct conversion translate directly to lower costs in many ways: fewer components means that an original-equipment manufacturer (OEM) spends less on bill of materials (BOM) and less for inserting the components (about a penny per insertion). The time to assemble a phone is reduced, increasing factory throughput; and the improved manufacturability of the phone (less to go wrong with fewer solder joints, etc.) increases reliability.




Figure 2. Photograph of a prototype Othello radio printed circuit board.

Because Othello radios can be so compact, they enable GSM radio technology to be incorporated in many products from which it has been excluded, such as very compact phones or PCMCIA cards. However, the real power of direct conversion will be seen when versatile third-generation phones are designed to handle multiple standards. With direct-conversion, hardware channel-selection filters will be unnecessary, because channel selection is performed in the digital signal-processing section, which can be programmed to handle multiple standards. Contrast this with the superheterodyne architecture; the multiple radio circuits required to handle the different standards (because each will require different channel-selection filters) will all have to be crowded into a small space. With direct conversion, the same radio chain could in concept be used for several different standards, bandwidths, and modulation types. Thus, Web-browsing and voice services could, in concept, occur over the GSM network using the same radio in the handset.

ANALOG DEVICES AND GSM

For the future, the Othello radio is only the first of a family of direct-conversion receiver solutions from Analog Devices. More are in the works. But this technology is well-grounded in nearly a decade of product designs for the GSM industry. ADI chips can be found in millions of GSM handsets in use worldwide. ▀

DSPs Enhance Flexible Third-Generation Base-Station Design

by Robert B. De Robertis and Rasekh Rifaat¹

Designers are working on the third generation of cellular infrastructure equipment. This new equipment will give the service provider higher user capacity (translates to more revenue), but more importantly, the consumer gets access to high-speed data services, including wireless Internet access and wireless video transmission. With target data rates of 2 megabits per second, standards for this generation are in the process of being completed by the 3rd Generation Partnership Project (3GPP). Ability to deliver flexible systems is the critical key for equipment manufacturers to win in this market. Today's proposals talk of achieving 2 MB/s data rates, but what about the 10-MB/s data rates of tomorrow? Flexible designs will enable the systems that are deployed to grow with the demands of the consumer and the needs of the service provider.

Such flexibility in system design is of financial importance to service providers because it extends the time during which their investments in equipment are paying off. Over time, systems that are programmable and scalable with minimal board replacement reduce the total capital outlay for new services. In turn, equipment manufacturers reduce their overall engineering costs in systems that are flexible and scalable to meet the future demands of the service providers—the demand for continually increased capabilities at lower cost.

To the design engineer, flexibility means many things, including consideration of new technologies. The heart of the 3G systems is raw signal processing, both analog and digital. Technologies and techniques, such as direct intermediate-frequency (IF) sampling, direct digital down conversion, digital signal processing, and re-configurable logic, enable more flexible base station design options today than were available for second-generation (2G) systems. Using them, the design engineer can create infrastructure equipment that has the flexibility to support the needs of service providers, and has the performance and throughput to be scalable as infrastructure demand grows.

DSPs ARE ABUNDANT WITHIN THE 3G SYSTEM

Key base-station areas that require high-performance DSPs include:

- Antenna Arrays with Adaptive Digital Beam-Forming (BS)
- Power Control (BS)
- Voice Processing (BSC: Base-Station Control)
- Baseband Modem (BTS: Base Transceiver Station)

The algorithms employed in these functional blocks are MAC-intensive (i.e., they employ many steps of *multiply-and-accumulate*). MAC-intensive functions for 3G include FIR, correlation, and

equalizer functions. The more rapidly these algorithms are performed, the better the quality and performance of a base station

The choice of a DSP to obtain the required computation speed is not a straightforward matter of specifying the highest clock speed. Architecture and instruction sets greatly affect the speed of algorithm execution. "MIPS" (millions of instructions per second) is also not a valid measure, since each manufacturer counts instructions differently. A highly useful recommended measure, more closely related to algorithm execution, is the peak *million-multiply-accumulates-per-second* (MMACS). This calculation is the product of the clock speed and the number of MACs the DSP is capable of executing per clock cycle².

Another aspect to consider is the class of DSP architecture employed. Two recently introduced new classes to consider are: *very long instruction word* (VLIW) and *static superscalar*.

VLIW attempts to reduce cost and increase execution speed by reducing hardware complexity. The sequencing mechanism in VLIW relies on an instruction format wherein every single execution unit in the chip is under direct programmer or compiler control. Unfortunately, VLIW has little or no hardware support for maintaining the integrity of data dependencies or avoiding scheduling hazards associated with real-time processing. In VLIW, all operation latencies in a particular implementation are fully exposed to software. The TMS320C6x series from Texas Instruments is an example of a VLIW architecture.

Static superscalar architectures enforce a consistent and functionally well-defined programming model, and the schedule is determined prior to run time. It incorporates static scheduling techniques like those found in VLIW, but it retains many superscalar and RISC attributes, enabling real time systems. Consequently, code can be written directly in assembly without requiring sophisticated timing prediction. The TigerSHARC™ DSP from Analog Devices is an example of a *static superscalar* architecture.

Antenna Arrays with Adaptive Digital Beam Forming

Digital beam-forming algorithms are designed to target source locations in a noisy environment. They rapidly compare responses of several spatially deployed antennas; the result of the computation is a signal that is believed to have originated from the target direction. Basically, they compute a correlation function that compares the signals and gives a measure of how close the desired and received signals are. Due to the many factors involved in the algorithm, and their wide dynamic range, floating point multiply-accumulate operations are used almost exclusively to minimize roundoff errors.

The target is mobile, and could be moving at a significant speed, this adds another dimension of complexity to the computation. *Adaptive beam-forming* makes use of additional information to continually track the mobile target. Beam-forming in 3G systems may be integrated with the *Rake* receiver, where the signal is operated on to combat fading and multipath effects. For these algorithms, the TigerSHARC's rapid performance of floating-point computations makes it an excellent fit.

TigerSHARC is a trademark of Analog Devices, Inc.

¹Much of the material in this article first appeared in *Wireless System Design*, published by Penton Media, Inc., October, 1999, and in their Web version, <http://www.wsdmag.com/>, November, 1999.

²D. Efsthathiou et al., "Recent Developments in Enabling Technologies for Software Defined Radio," *IEEE Communications magazine*, Aug., 1999.

³RT. Ojanpera, *Wideband CDMA for Third Generation Mobile Communications*, Boston-London: Artech House, 1998.

Power Control

In the code-division multiple-access (CDMA) systems proposed for 3G, base-station-initiated power control of remote-unit transmitters (uplink) is critical to compensate for fast fading, peaks in transmission power, and to avoid near-far problems³. This is necessary to reduce inter-cell interference. The computations required for power control are multiply-accumulate intensive, requiring high performance digital signal processing to meet delay time requirement in 3G systems.

Base stations may also implement the feedback mode transmit diversity (FMTD) algorithm, which is a power control/beam forming application that uses multiple antenna transmissions with varying weights. Again, the computation is multiply-accumulate intensive, similarly to the rake receiver. For such applications, the ADSP-21065L SHARC is a processor of choice.

Voice Processing

DSPs are the traditional choice for speech processing within the cellular system. The phone user's opinion of the quality of the system is directly dependent on the performance of the speech coder, and this has a strong influence on the channel density. Several speech coders are in use today in current 2G systems and must be supported in 3G systems. (See Table 1). Although lower codec bit rates increase equipment capacity, they worsen the speech quality. The critical DSP characteristics for high-quality voice processing combine large on-chip RAM and high processing capacity to support fast context switching and high channel density. The ADSP-21mod980, with its 8 DSP cores, capable of 600 MMACS (Million MACs per second), is the ideal candidate for this portion of the signal chain.

Table 1. Cellular Speech Codecs

Speech Codec	Bit Rate (Kbps)	Standard
QCELP	8	IS-95
EVRC	Variable	IS-95
ACELP	13	IS-95
VSELP	8	IS-136
GSM FR	13	GSM
GSM EFR	12.2	GSM
AMR	Variable	3G
GSM HR	5.6	GSM
JVSELP	8	PDC

Baseband Modem

The 3G standard is expected to be an essential factor that enables applications involving the transmission of wideband signals. Accordingly, the baseband modem (BTS) must be designed and implemented with the ability to intermix high bandwidth applications and low bandwidth voice and paging. In the downlink, the base transceiver station packages parallel transport-block streams into physical channels; and in the uplink, it recreates the transport blocks from the baseband signal.

Figure 2 shows a typical baseband modem section of the 3G base station for both uplink and downlink configurations. During downlink, error-coding schemes are first applied to the transport block. Then the blocks are reordered and recombined with other channels before being sent off to the radio. For the uplink, the rake receiver is first used to sort out multipath effects and possibly to combine the data from several antennas. The blocks are then restored to their original order and channels before forward error correction is applied.

In the next section, the partitioning of the baseband modem provides insight as to where a designer might choose to use a DSP. An optimum must be sought between minimizing the performance cost and maximizing the flexibility of the system to handle future design iterations.

BASEBAND MODEM PARTITIONING

Careful Partitioning for Maximum Flexibility

When deciding how to partition the modem, the nature of the algorithms and data rates become key factors in deciding what should be processed with an ASIC and what should be performed in the DSP. The next section describes some of these algorithms in detail, and explains the tradeoffs. This guideline depends upon the cost to process N channels of a specified bandwidth. For maximum flexibility, the entire structure could be implemented using a cluster of DSPs. On the other hand, a fixed specification may be most-economically implemented in an ASIC. When evaluating the most suitable approach, the flexibility criterion demands that an engineer consider how the design supports:

- quickly upgrading the parts of a system to newer technology,
- scaling the system to improve performance,
- product differentiation through the addition of new features.

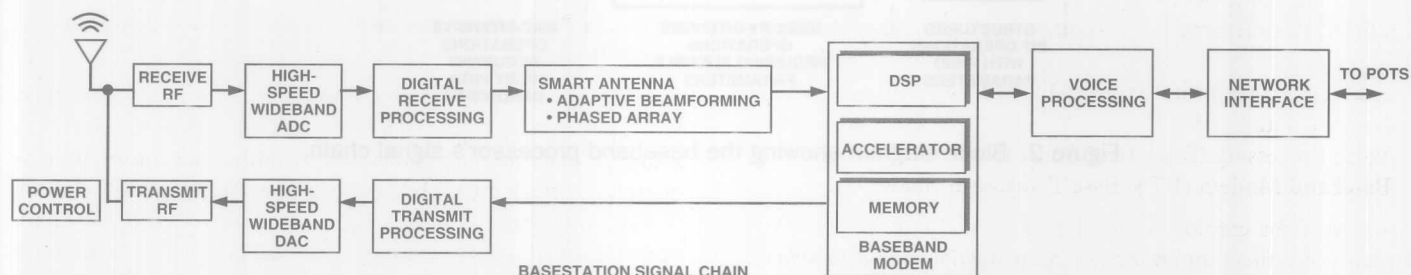


Figure 1. Where to find DSP technology in 3G systems.

Rake, Channel-Encoding/Decoding Hardware-Software Tradeoffs

Figure 2 shows different classifications for the different parts of the modem. Each of the blocks performs different types of computations. An overview is needed in order to see exactly where DSP is more appropriate than other alternatives.

The interleaving, channel segmenting, and rate matching are I/O-intensive operations, which combine data from several sources and reorganize data to minimize the effects of errors. Because of the variability of the parameters, data-rates, and memory-referencing, these functions are ideally suited to DSP for manipulation; they would be difficult to implement cost-effectively in an ASIC.

The error-coding and -correction algorithms involve significant bit manipulations that—properly implemented—can be implemented in the DSP. The error-correction algorithms also represent an area of the modem that can provide equipment-manufacturer differentiation. The encoding standards have been fixed, while decoding is left to implementers to design using their own intellectual property. Companies that have a strong ASIC capability might choose a hard-wired design, while others that are strong in programming and desire flexibility will choose the DSP approach. ADI's TigerSHARC DSP provides all the processing capacity to enable a single high speed 3G data channel.

3G systems, employing spread spectrum communications, will utilize CDMA spreading codes in order to provide greater use of available bandwidth. The spreading and despreading algorithms are multiply-accumulate intensive, but at extremely high data rates. The rake receiver takes its name from the fact that its diagram resembles a garden rake. Each finger tries to correlate the incoming data with the expected spreading code. As a result, the rake receiver needs to be able to process K times

the determined bandwidth, where K is the number of fingers in the rake. In addition to that, the receiver must operate at the frequency set by the spreading codes.

DSP technology today doesn't cost-effectively support the bandwidth required for the spreading and despreading in 3G systems. However, production systems will not be manufactured and installed for several years, so there is every incentive for this situation to change. A clear examination of the upgrade path and projected performance of DSPs may show that, by the time these systems are deployed, these functions can indeed be handled cost-effectively.

Glueless Homogeneous and Heterogeneous Multiprocessing

Regardless of the technology used to implement each section of the baseband modem, a significant amount of data must be moved around the system. In considering the design, components (or groups of components) that support high-bandwidth communications must be used. The TigerSHARC DSP provides several options for high-speed communication, including on-chip DMA (direct memory-access) and SDRAM support, along with dedicated user-programmable link ports. In multiprocessing designs, a high speed cluster bus can be used to connect as many as eight TigerSHARC DSPs without additional logic.

CONCLUSION

Designers of 3G base stations will make use of the DSPs in order to achieve the high performance and flexibility needed for tomorrow's voice and data applications. Flexibility at all levels will drive the need for scalable technologies, such as static super scalar architectures and glueless interconnection of system components. Effective embodiment of these design principles will fulfill the promise of 3G to provide the foundations of the kind of wireless infrastructure necessary for tomorrow's killer applications. ▀

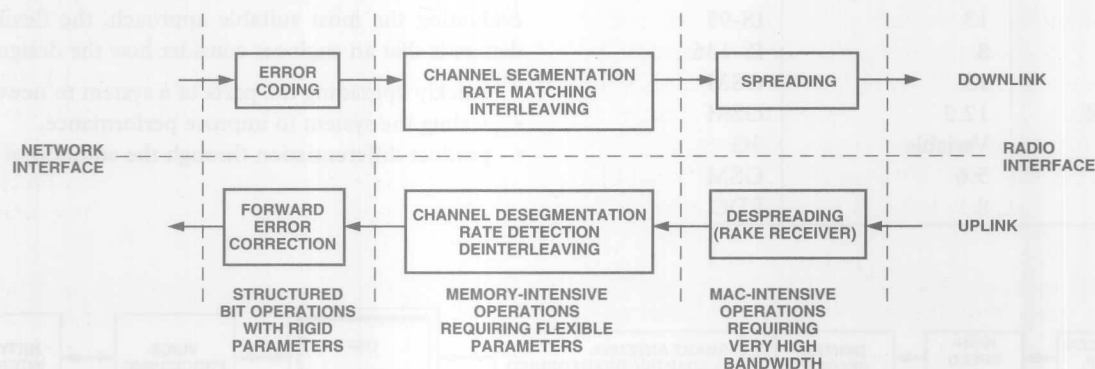


Figure 2. Block diagram showing the baseband processor's signal chain.

Phase-Locked Loops for High-Frequency Receivers and Transmitters—Part 1

by Mark Curtin and Paul O'Brien

This 3-part series of articles is intended to give a comprehensive overview of the use of PLLs (phase-locked loops) in both wired and wireless communication systems.

In this first part, the emphasis is on the introductory concepts of PLLs. The basic PLL architecture and principle of operation is described. We will also give an example of where PLLs are used in communication systems. We will finish the first installment by showing a practical PLL circuit using the ADF4111 Frequency Synthesizer and the VCO190-902T Voltage-Controlled Oscillator.

In the second part, we will examine in detail the critical specifications associated with PLLs: phase noise, reference spurs and output leakage current. What causes these and how can they be minimized? What effect do they have on system performance?

The final installment will contain a detailed description of the blocks that go to make up a PLL synthesizer and the architecture of an Analog Devices synthesizer. There will also be a summary of synthesizers and VCOs currently available on the market, with a list of ADI's current offerings.

PLL BASICS

A phase-locked loop is a feedback system combining a voltage-controlled oscillator and a phase comparator so connected that the oscillator maintains a constant phase angle relative to a reference signal. Phase-locked loops can be used, for example, to generate stable output frequency signals from a fixed low-frequency signal. The first phase-locked loops were implemented in the early 1930s by a French engineer, de Bellescize. However, they only found broad acceptance in the marketplace when integrated PLLs became available as relatively low-cost components in the mid-1960s.

The phase locked loop can be analyzed in general as a negative-feedback system with a forward gain term and a feedback term.

A simple block diagram of a voltage-based negative-feedback system is shown in Figure 1.




Figure 1. Standard negative-feedback control system model.

In a phase-locked loop, the error signal from the phase comparator is proportional to the relative phase of the input and feedback signals. The average output of the phase detector will be constant when the input and feedback signals are the same frequency. The usual equations for a negative-feedback system apply.

$$\text{Forward Gain} = G(s), [s = j\omega = j2\pi f]$$

$$\text{Loop Gain} = G(s) \times H(s)$$

$$\text{Closed-Loop Gain} = \frac{G(s)}{1 + G(s)H(s)}$$

Because of the integration in the loop, at low frequencies the steady state gain, $G(s)$, is high and

$$V_O/V_I, \text{ Closed-Loop Gain} = \frac{1}{H}$$

The components of a PLL that contribute to the loop gain include:

1. The *phase detector* (PD) and *charge pump* (CP).
2. The *loop filter*, with a transfer function of $Z(s)$
3. The *voltage-controlled oscillator* (VCO), with a sensitivity of K_V/s
4. The *feedback divider*, $1/N$




Figure 2. Basic phase-locked-loop model.

If a linear element like a four-quadrant multiplier is used as the phase detector, and the loop filter and VCO are also analog elements, this is called an analog, or *linear PLL* (LPLL).

If a *digital* phase detector (EXOR gate or J-K flip flop) is used, and everything else stays the same, the system is called a *digital PLL* (DPLL).

If the PLL is built exclusively from digital blocks, without any passive components or linear elements, it becomes an *all-digital PLL* (ADPLL).

Finally, with information in digital form, and the availability of sufficiently fast processing, it is also possible to develop PLLs in the software domain. The PLL function is performed by software and runs on a DSP. This is called a *software PLL* (SPLL).

Referring to Figure 2, a system for using a PLL to generate higher frequencies than the input, the VCO oscillates at an angular frequency of ω_O . A portion of this signal is fed back to the error detector, via a frequency divider with a ratio $1/N$. This divided-down frequency is fed to one input of the error detector. The other input in this example is a fixed reference signal. The error detector compares the signals at both inputs. When the two signal inputs are equal in frequency, the error will be constant and the loop is said to be in a "locked" condition. If we simply look at the error signal, the following equations may be developed.

$$e(s) = \Phi_{REF} - \frac{\Phi_O}{N}$$

$$\frac{de(s)}{dt} = F_{REF} - \frac{F_O}{N}$$

When

$$e(s) = \text{constant}, \frac{F_O}{N} = F_{REF}$$

Thus

$$F_O = N F_{REF}$$

In commercial PLLs, the phase detector and charge pump together form the error detector block. When $F_O \neq N F_{REF}$, the error detector will output source/sink current pulses to the low-pass loop filter. This smooths the current pulses into a voltage which in turn drives the VCO. The VCO frequency will then increase or decrease as necessary, by $K_V DV$, where K_V is the VCO sensitivity in MHz/Volt and DV is the change in VCO input voltage. This will continue until $e(s)$ is zero and the loop is locked. The charge pump and VCO thus serves as an integrator, seeking to increase or decrease its output frequency to the value required so as to restore its input (from the phase detector) to zero.




Figure 3. VCO transfer function.

The overall transfer function (CLG or Closed-Loop Gain) of the PLL can be expressed simply by using the CLG expression for a negative feedback system as given above.

$$\frac{F_O}{F_{REF}} = \frac{\text{Forward Gain}}{1 + \text{Loop Gain}}$$

$$\text{Forward Gain, } G = \frac{K_D K_V Z(s)}{s}$$

$$\text{Loop Gain, } GH = \frac{K_D K_V Z(s)}{Ns}$$

Table 2. Channel Numbering for GSM900 and DCS1800 Base Station Systems

	Rx		Tx
PGSM900	$F_l(n) = 890 + 0.2 \times (n)$	$1 \leq n \leq 124$	$F_u(n) = F_l(n) + 45$
EGSM900	$F_l(n) = 890 + 0.2 \times (n)$ $F_l(n) = 890 + 0.2 \times (n - 1024)$	$0 \leq n \leq 124$ $975 \leq n \leq 1023$	$F_u(n) = F_l(n) + 45$
DCS1800	$F_l(n) = 1710.2 + 0.2 \times (n - 512)$	$512 \leq n \leq 885$	$F_u(n) = F_l(n) + 95$

When GH is much greater than 1, we can say that the closed loop transfer function for the PLL system is N and so

$$F_{OUT} = N \times F_{REF}$$

The loop filter is a low-pass type, typically with one pole and one zero. The transient response of the loop depends on:

1. the magnitude of the pole/zero,
2. the charge pump magnitude,
3. the VCO sensitivity,
4. the feedback factor, N .

All of the above must be taken into account when designing the loop filter. In addition, the filter must be designed to be stable (usually a phase margin of $\pi/4$ is recommended). The 3-dB cutoff frequency of the response is usually called the loop bandwidth, B_w . Large loop bandwidths result in very fast transient response. However, this is not always advantageous, as we shall see in Part 2, since there is a tradeoff between fast transient response and reference spur attenuation.

PLL APPLICATIONS TO FREQUENCY UPSCALING

The phase-locked loop allows stable high frequencies to be generated from a low-frequency reference. Any system that requires stable high frequency tuning can benefit from the PLL technique. Examples of these applications include wireless base stations, wireless handsets, pagers, CATV systems, clock-recovery and -generation systems. A good example of a PLL application is a GSM handset or base station. Figure 4 shows the receive section of a GSM base station.

In the GSM system, there are 124 channels (8 users per channel) of 200-kHz width in the RF band. The total bandwidth occupied is 24.8 MHz, which must be scanned for activity. The handset has a transmit (Tx) range of 880 MHz to 915 MHz and a receive (Rx) range of 925 MHz to 960 MHz. Conversely, the base station has a Tx range of 925 MHz to 960 MHz and an Rx range of 880 MHz to 915 MHz. For this example, we will consider just the base station transmit and receive sections. The frequency bands for GSM900 and DCS1800 Base Station Systems are shown in Table 1. Table 2 shows the channel numbers for the carrier frequencies (RF channels) within the frequency bands of Table 1. $F_l(n)$ is the center frequency of the RF channel in the lower band (Rx) and $F_u(n)$ is the corresponding frequency in the upper band (Tx).

Table 1. Frequency Bands for GSM900 and DCS1800 Base Station Systems

	Tx	Rx
P-GSM900	935 to 960 MHz	890 to 915 MHz
DCS1800	1805 to 1880 MHz	1710 to 1785 MHz
E-GSM900	925 to 960 MHz	880 to 915 MHz

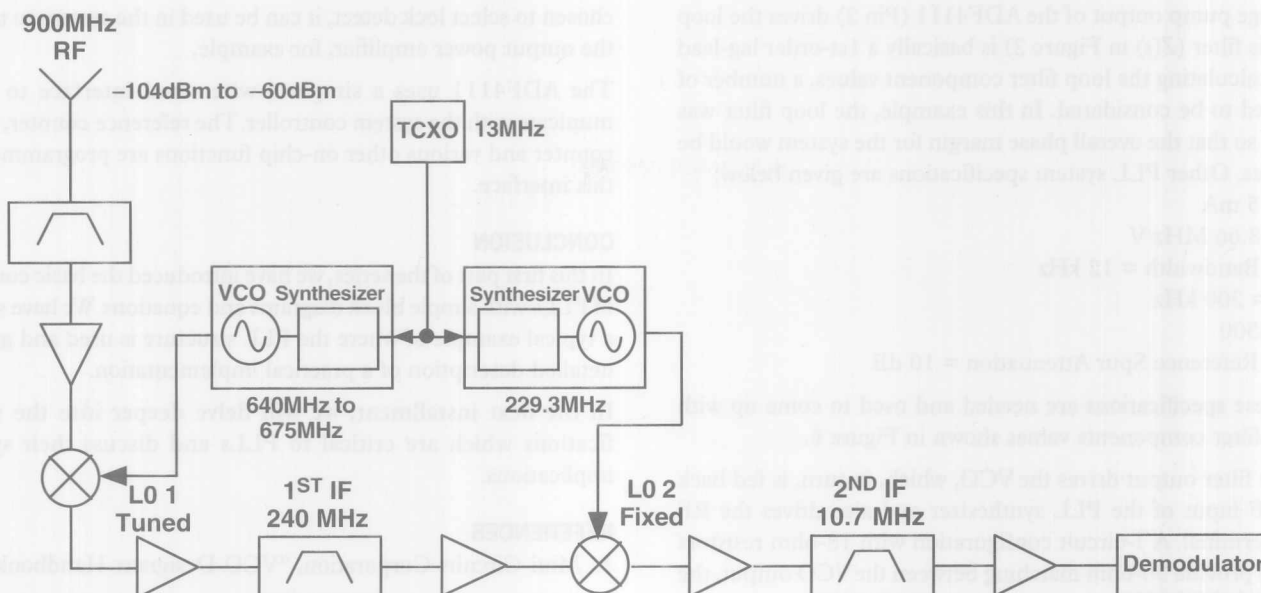


Figure 4. Signal chain for GSM base-station receiver.

The 900-MHz RF input is filtered, amplified and applied to the first stage mixer. The other mixer input is driven from a tuned local oscillator (LO). This must scan the input frequency range to search for activity on any of the channels. The actual implementation of the LO is by means of the PLL technique already described. If the 1st intermediate-frequency (IF) stage is centered at 240 MHz, then the LO must have a range of 640 MHz to 675 MHz in order to cover the RF input band. When a 200-kHz reference frequency is chosen, it will be possible to sequence the VCO output through the full frequency range in steps of 200 kHz. For example, when an output frequency of 650 MHz is desired, N will have a value of 3250. This 650-MHz LO will effectively check the 890-MHz RF channel ($F_{RF} - F_{LO} = F_{IF}$ or $F_{RF} = F_{LO} + F_{IF}$). When N is incremented to 3251, the LO frequency will now be 650.2 MHz and the RF channel checked will be 890.2 MHz. This is shown graphically in Figure 5.




Figure 5. Testing frequencies for GSM base-station receiver.

It is worth noting that, in addition to the tunable RF LO, the receiver section also uses a fixed IF (in the example shown this is 240 MHz). Even though frequency tuning is not needed on this IF, the PLL technique is still used. The reason for this is that it is an affordable way of using the stable system reference frequency to produce the high frequency IF signal. Several synthesizer manufacturers recognize this fact by offering dual versions of the devices: one operating at the high RF frequency (>800 MHz) and one operating at the lower IF frequency (500 MHz or less).

On the transmit side of the GSM system, similar requirements exist. However, it is more common to go directly from baseband to the final RF in the Transmit section; this means that the typical T_x VCO for a base station has a range of 925 MHz to 960 MHz (RF band for the Transmit section).

CIRCUIT EXAMPLE

Figure 6 shows an actual implementation of the local oscillator for the transmit section of a GSM handset. We are assuming direct baseband to RF up-conversion. This circuit uses the new ADF4111 PLL Frequency Synthesizer from ADI and the VCO190-902T Voltage Controlled Oscillator from Vari-L Corporation (<http://www.vari-l.com/>).

The reference input signal is applied to the circuit at $F_{REF_{IN}}$ and is terminated in 50 Ω . This reference input frequency is typically 13 MHz in a GSM system. In order to have a channel spacing of 200 kHz (the GSM standard), the reference input must be divided by 65, using the on-chip reference divider of the ADF4111.

The ADF4111 is an integer-N PLL frequency synthesizer, capable of operating up to an RF frequency of 1.2 GHz. In this integer-N type of synthesizer, N can be programmed from 96 to 262,000 in discrete integer steps. In the case of the handset transmitter, where an output range of 880 MHz to 915 MHz is needed, and where the internal reference frequency is 200 kHz, the desired N values will range from 4400 to 4575.

The charge pump output of the ADF4111 (Pin 2) drives the loop filter. This filter ($Z(s)$ in Figure 2) is basically a 1st-order lag-lead type. In calculating the loop filter component values, a number of items need to be considered. In this example, the loop filter was designed so that the overall phase margin for the system would be 45 degrees. Other PLL system specifications are given below:

$$K_D = 5 \text{ mA}$$

$$K_V = 8.66 \text{ MHz/V}$$

$$\text{Loop Bandwidth} = 12 \text{ kHz}$$

$$F_{\text{REF}} = 200 \text{ kHz}$$

$$N = 4500$$

$$\text{Extra Reference Spur Attenuation} = 10 \text{ dB}$$

All of these specifications are needed and used to come up with the loop filter components values shown in Figure 6.

The loop filter output drives the VCO, which, in turn, is fed back to the RF input of the PLL synthesizer and also drives the RF Output terminal. A T-circuit configuration with 18-ohm resistors is used to provide 50-ohm matching between the VCO output, the RF output and the RF_{IN} terminal of the ADF4111.

In a PLL system, it is important to know when the system is in lock. In Figure 6, this is accomplished by using the MUXOUT signal from the ADF4111. The MUXOUT pin can be programmed to monitor various internal signals in the synthesizer. One of these is the LD or *lock-detect* signal. When MUXOUT is

chosen to select lock detect, it can be used in the system to trigger the output power amplifier, for example.

The ADF4111 uses a simple 4-wire serial interface to communicate with the system controller. The reference counter, the N counter and various other on-chip functions are programmed via this interface.

CONCLUSION

In this first part of the series, we have introduced the basic concepts of PLLs with simple block diagrams and equations. We have shown a typical example of where the PLL structure is used and given a detailed description of a practical implementation.

In the next installment, we will delve deeper into the specifications which are critical to PLLs and discuss their system implications.

REFERENCES

1. Mini-Circuits Corporation, "VCO Designers Handbook."
2. L.W. Couch, "Digital and Analog Communications Systems" Macmillan Publishing Company, New York.
3. P. Vizmuller, "RF Design Guide," Artech House.
4. R.L. Best, "Phase Locked Loops: Design, Simulation and Applications," 3rd Edition, McGraw Hill. ▣

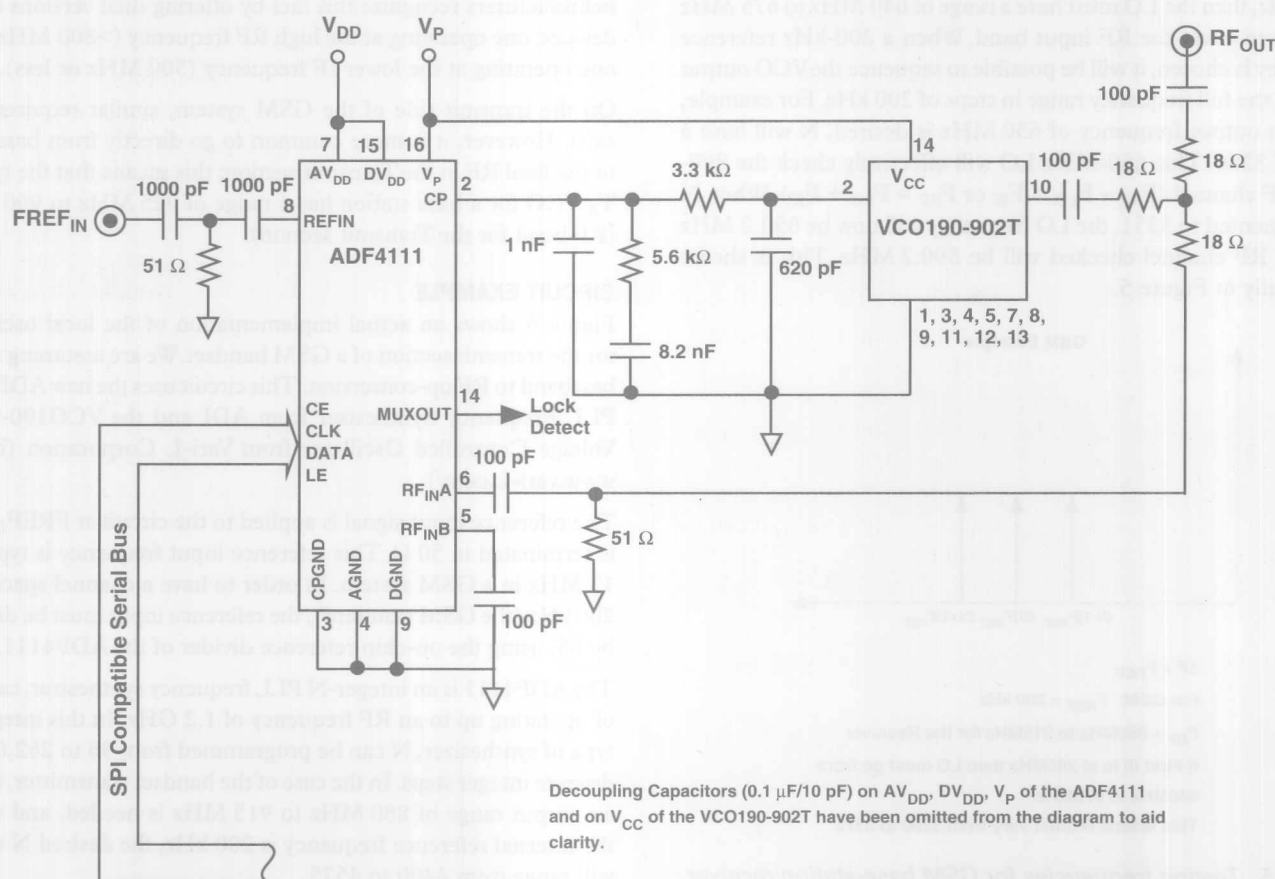


Figure 6. Transmitter local oscillator for GSM handset.

Phase-Locked Loops for High-Frequency Receivers and Transmitters—Part 2

by Mike Curtin and Paul O'Brien

The *first part* of this series of articles introduced the basic concepts of *phase-locked loops* (PLLs). The PLL architecture and principle of operation was described and accompanied by an example of where a PLL might be used in a communication system.

In this second part, we will focus on a detailed examination of two critical specifications associated with PLLs: *phase noise* and *reference spurs*. What causes them and how can they be minimized? The discussion will include measurement techniques and the effect of these errors on system performance. We will also consider *output leakage current*, with an example showing its significance in open-loop modulation schemes.

Noise in Oscillator Systems

In any oscillator design, frequency stability is of critical importance. We are interested in both long-term and short-term stability. *Long-term* frequency stability is concerned with how the output signal varies over a long period of time (hours, days or months). It is usually specified as the ratio, $\Delta f/f$ for a given period of time, expressed as a percentage or in dB.

Short-term stability, on the other hand, is concerned with variations that occur over a period of seconds or less. These variations can be random or periodic. A spectrum analyzer can be used to examine the short-term stability of a signal. Figure 1 shows a typical spectrum, with random and discrete frequency components causing a broad skirt and spurious peaks.




Figure 1. Short-term stability in oscillators.

The discrete spurious components could be caused by known clock frequencies in the signal source, power line interference, and mixer products. The broadening caused by random noise fluctuation is due to *phase noise*. It can be the result of thermal noise, shot noise and/or flicker noise in active and passive devices.

Phase Noise in Voltage-Controlled Oscillators

Before we look at phase noise in a PLL system, it is worth considering the phase noise in a *voltage-controlled oscillator* (VCO). An ideal VCO would have no phase noise. Its output as seen on a spectrum analyzer would be a single spectral line. In practice, of course, this is not the case. There will be jitter on the output, and a spectrum analyzer would show phase noise. To help understand phase noise, consider a phasor representation, such as that shown in Figure 2.




Figure 2. Phasor representation of phase noise.

A signal of angular velocity ω_0 and peak amplitude V_{SPK} is shown. Superimposed on this is an error signal of angular velocity ω_m . $\Delta\theta_{rms}$ represents the rms value of the phase fluctuations and is expressed in rms degrees.

In many radio systems, an overall integrated phase error specification must be met. This overall phase error is made up of the PLL phase error, the modulator phase error and the phase error due to base band components. In GSM, for example, the total allowed is 5 degrees rms.

Leeson's Equation

Leeson (see Reference 6) developed an equation to describe the different noise components in a VCO.

$$L_{PM} \approx 10 \log \left[\frac{FkT}{A} \frac{1}{8Q_L^2} \left(\frac{f_0}{f_m} \right)^2 \right] \quad (1)$$

where:

L_{PM} is single-sideband phase noise density (dBc/Hz)

F is the device noise factor at operating power level A (linear)

k is Boltzmann's constant, 1.38×10^{-23} J/K

T is temperature (K)

A is oscillator output power (W)

Q_L is loaded Q (dimensionless)

f_0 is the oscillator carrier frequency

f_m is the frequency offset from the carrier

For Leeson's equation to be valid, the following must be true:

- f_m , the offset frequency from the carrier, is greater than the $1/f$ flicker corner frequency;
- the noise factor at the operating power level is known;
- the device operation is linear;
- Q includes the effects of component losses, device loading and buffer loading;
- a single resonator is used in the oscillator.

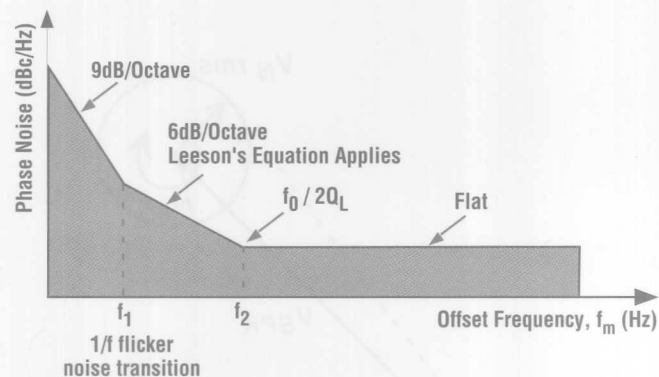


Figure 3. Phase noise in a VCO vs. frequency offset.

Leeson's equation only applies in the knee region between the break (f_1) to the transition from the "1/f" (more generally $1/f^\gamma$) flicker noise frequency to a frequency beyond which amplified white noise dominates (f_2). This is shown in Figure 3 [$\gamma = 3$]. f_1 should be as low as possible; typically, it is less than 1 kHz, while f_2 is in the region of a few MHz. High-performance oscillators require devices specially selected for low 1/f transition frequency. Some guidelines to minimizing the phase noise in VCOs are:

1. Keep the tuning voltage of the varactor sufficiently high (typically between 3 and 3.8 V)
2. Use filtering on the dc voltage supply.
3. Keep the inductor Q as high as possible. Typical off-the-shelf coils provide a Q of between 50 and 60.
4. Choose an active device that has minimal noise figure as well as low flicker frequency. The flicker noise can be reduced by the use of feedback elements.
5. Most active device exhibit a broad U-shaped noise-figure-vs.-bias-current curve. Use this information to choose the optimal operating bias current for the device.
6. Maximize the average power at the tank circuit output.
7. When buffering the VCO, use devices with the lowest possible noise figure.

Closing The Loop

Having looked at phase noise in a free-running VCO and considered how it can be minimized, we will now consider the effect of closing the loop (see *Part 1* of the series) on phase noise.




Figure 4. PLL-phase-noise contributors.

Figure 4 shows the main phase noise contributors in a PLL. The system transfer function may be described by the following equations.

$$\text{Closed Loop Gain} = \frac{G}{1 + GH} \quad (2)$$

$$G = \frac{K_d \times K_v \times Z(s)}{s} \quad (3)$$

$$H = \frac{1}{N} \quad (4)$$

$$\text{Closed Loop Gain} = \frac{\frac{K_d \times K_v \times Z(s)}{s}}{\frac{K_d \times K_v \times Z(s)}{s} + N} \quad (5)$$

For the discussion that follows, we will define S_{REF} as the noise that appears on the reference input to the phase detector. It is dependent on the reference divider circuitry and the spectral purity of the main reference signal. S_N is the noise due to the feedback divider appearing at the frequency input to the phase detector. S_{CP} is the noise due to the phase detector (depending on its implementation). And S_{VCO} is the phase noise of the VCO as described by equations developed earlier.

The overall phase noise performance at the output depends on the terms described above. All the effects at the output are added in an rms fashion to give the total noise of the system. Thus:

$$S_{TOT}^2 = X^2 + Y^2 + Z^2 \quad (6)$$

where:

S_{TOT}^2 is the total phase noise power at the output

X^2 is the noise power at the output due to S_N and S_{REF} .

Y^2 is the noise power at the output due to S_{CP} .

Z^2 is the noise power at the output due to S_{VCO} .

The noise terms at the PD inputs, S_{REF} and S_N , will be operated on in the same fashion as F_{REF} and will be multiplied by the closed loop gain of the system.

$$X^2 = \left(S_{REF}^2 + S_N^2 \right) \times \left(\frac{G}{1 + GH} \right)^2 \quad (7)$$

At low frequencies, inside the loop bandwidth,

$$GH \gg 1 \text{ and } X^2 = \left(S_{REF}^2 + S_N^2 \right) \times N^2 \quad (8)$$

At high frequencies, outside the loop bandwidth,

$$G \ll 1 \text{ and } X^2 \Rightarrow 0 \quad (9)$$

The overall output noise contribution due to the phase detector noise, S_{CP} , can be calculated by referencing S_{CP} back to the input of the PFD. The equivalent noise at the PD input is S_{CP}/K_d . This is then multiplied by the closed-loop gain:

$$Y^2 = S_{CP}^2 \times \left(\frac{1}{K_d} \right)^2 \times \left(\frac{G}{1 + GH} \right)^2 \quad (10)$$

Finally, the contribution of the VCO noise, S_{VCO} , to the output phase noise is calculated in a similar manner. The forward gain this time is simply 1. Therefore its contribution to the output noise is:

$$Z^2 = S_{VCO}^2 \times \left(\frac{1}{1 + GH} \right)^2 \quad (11)$$

G , the forward loop gain of the closed loop response, is usually a low pass function; it is very large at low frequencies and small at high frequencies. H is a constant, $1/N$. The denominator of the above expression is therefore low pass, so S_{VCO} is actually high-pass filtered by the closed loop.

A similar description of the noise contributors in a PLL/VCO can be found in Reference 1. Recall that the closed-loop response is a low-pass filter with a 3-dB cutoff frequency, B_W , denoted the *loop bandwidth*. For frequency offsets at the output less than B_W , the dominant terms in the output phase noise response are X and Y , the noise terms due to reference noise, N (counter noise), and charge pump noise. Keeping S_N and S_{REF} to a minimum, keeping K_d large and keeping N small will thus minimize the phase noise inside the loop bandwidth, B_W . Because N programs the output frequency, it is not generally available as a factor in noise reduction.

For frequency offsets much greater than B_W , the dominant noise term is that due to the VCO, S_{VCO} . This is due to the high pass filtering of the VCO phase noise by the loop. A small value of B_W would be desirable as it would minimize the total integrated output noise (phase error). However a small B_W results in a slow transient response and increased contribution from the VCO phase noise inside the loop bandwidth. The loop bandwidth calculation therefore must trade off transient response and total output integrated phase noise.

To show the effect of closing the loop on a PLL, Figure 5 shows an overlay of the output of a free-running VCO and the output of a VCO as part of a PLL. Note that the in-band noise of the PLL has been attenuated compared to that of the free-running VCO.




Figure 5. Phase noise on a free-running VCO and a PLL-connected VCO.

Phase Noise Measurement

One of the most common ways of measuring phase noise is with a high frequency spectrum analyzer. Figure 6 is a typical example of what would be seen.




Figure 6. Phase noise definition.

With the spectrum analyzer we can measure the spectral density of phase fluctuations per unit bandwidth. VCO phase noise is best described in the frequency domain where the spectral density is characterized by measuring the noise sidebands on either side of the output signal center frequency. Phase noise power is specified in decibels relative to the carrier (dBc/Hz) at a given frequency offset from the carrier. The following equation describes this SSB phase noise (dBc/Hz).

$$S_C(f) = 10 \log \frac{P_s}{P_{SSB}} \quad (12)$$

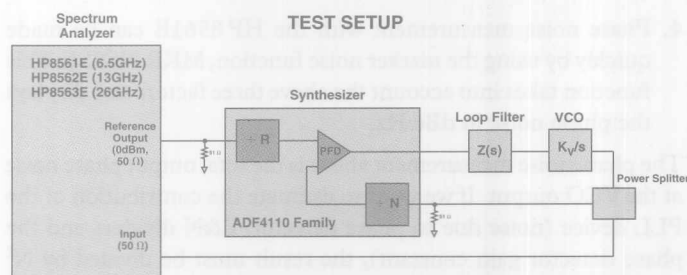


Figure 7. Measuring phase noise with a spectrum analyzer.

The 10-MHz, 0-dBm reference oscillator, available on the spectrum analyzer's rear-panel connector, has excellent phase noise performance. The R divider, N divider, and the phase detector are part of ADF4112 frequency synthesizer. These dividers are programmed serially under the control of a PC. The frequency and phase noise performance are observed on the spectrum analyzer.

Figure 8 illustrates a typical phase noise plot of a PLL synthesizer using an ADF4112 PLL with a Murata VCO, MQE520-1880. The frequency and phase noise were measured in a 5-kHz span. The reference frequency used was $f_{REF} = 200$ kHz ($R = 50$) and the output frequency was 1880 MHz ($N = 9400$). If this were an ideal-world PLL synthesizer, a single discrete tone would be displayed rising up above the spectrum analyzer's noise floor. What is displayed here is the tone, with the phase noise due to the loop components. The loop filter values were chosen to give a loop




Figure 8. Typical spectrum-analyzer output.

bandwidth of approximately 20 kHz. The flat part of the phase noise for frequency offsets less than the loop bandwidth is actually the phase noise as described by X^2 and Y^2 in the section "closing the loop" for cases where f is inside the loop bandwidth. It is specified at a 1-kHz offset. The value measured, the phase-noise power in a 1-Hz bandwidth, was -85.86 dBc/Hz. It is made up of the following:

1. Relative power in dBc between the carrier and the sideband noise at 1-kHz offset.
2. The spectrum analyzer displays the power for a certain resolution bandwidth (RBW). In the plot, a 10-Hz RBW is used. To represent this power in a 1-Hz bandwidth, $10\log(\text{RBW})$ must be subtracted from the value obtained from (1).
3. A correction factor, which takes into account the implementation of the RBW, the log display mode and detector characteristic, must be added to the result obtained in (2).
4. Phase noise measurement with the HP 8561E can be made quickly by using the marker noise function, MKR NOISE. This function takes into account the above three factors and displays the phase noise in dBc/Hz.

The phase noise measurement above is the total output phase noise at the VCO output. If we want to estimate the contribution of the PLL device (noise due to phase detector, R&N dividers and the phase detector gain constant), the result must be divided by N^2 (or $20 \times \log N$ be subtracted from the above result). This gives a phase-noise floor of $[-85.86 - 20 \times \log(9400)] = -165.3$ dBc/Hz.

Reference Spurs

In an integer-N PLL (where the output frequency is an integer multiple of the reference input), reference spurs are caused by the fact that the charge pump output is being continuously updated at the reference frequency rate. Consider




Figure 9. Basic PLL model.

again the basic model for the PLL which was discussed in Part 1 of this series. This is shown again in Figure 9.

When the PLL is in lock, the phase and frequency inputs to the PFD (f_{REF} and f_N) are essentially equal, and, in theory, one would expect that there to be no output from the PFD. However, this can create problems (to be discussed in Part 3 of this series), so the PFD is designed such that, in the locked condition, the current pulses from the charge pump will typically be as shown in Figure 10.




Figure 10. Output current pulses from the PFD charge pump.

Although these pulses have a very narrow width, the fact that they exist means that the dc voltage driving the VCO is modulated by a signal of frequency f_{REF} . This produces *reference spurs* in the RF output occurring at offset frequencies that are integer multiples of f_{REF} . A spectrum analyzer can be used to detect reference spurs. Simply increase the span to greater than twice the reference frequency. A typical plot is shown in Figure 11. In this case the reference frequency is 200 kHz and the diagram clearly shows reference spurs at ± 200 kHz from the RF output of 1880 MHz. The level of these spurs is -90 dB. If the span were increased to more than four times the reference frequency, we would also see the spurs at $(2 \times f_{\text{REF}})$.




Figure 11. Output spectrum showing reference spurs.

Charge Pump Leakage Current

When the CP output from the synthesizer is programmed to the high impedance state, there should, in theory, be no leakage current flowing. In practice, in some applications the level of leakage current will have an impact on overall system performance. For example, consider an application where a PLL is used in open-loop mode for frequency modulation—a simple and inexpensive way of implementing FM that also allows higher data rates than modulating in closed-loop mode. For FM, a closed-loop method works fine but the data rate is limited by the loop bandwidth.

A system that uses open-loop modulation is the European cordless telephone system, DECT. The output carrier frequencies are in a range of 1.77 GHz to 1.90 GHz and the data rate is high; 1.152 Mbps.

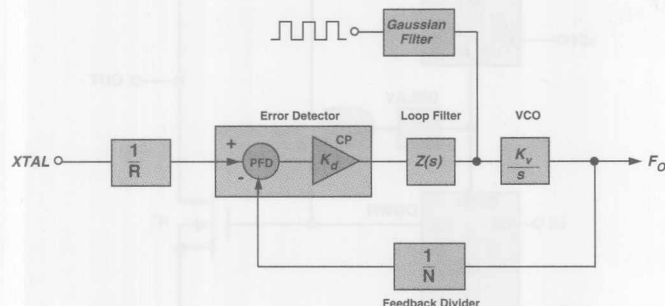


Figure 12. Block diagram of open-loop modulation.

A block diagram of open-loop modulation is shown in Figure 12. The principle of operation is as follows: The loop is initially closed to lock the RF output, $f_{OUT} = N f_{REF}$. The modulating signal is turned on and at first the modulation signal is simply the dc mean of the modulation. The loop is then opened, by putting the CP output of the synthesizer into high-impedance mode, and the modulation data is fed to the Gaussian filter. The modulating voltage then appears at the VCO where it is multiplied by K_V . When the data burst finishes, the loop is returned to the closed loop mode of operation.

As the VCO usually has a high sensitivity (typical figures are between 20 and 80 MHz/volt), any small voltage drift before the VCO will cause the output carrier frequency to drift. This voltage drift, and hence the system frequency drift, is directly dependent on the leakage current of the charge pump, CP, when in the high impedance state. This leakage will cause the loop capacitor to charge or discharge depending on the polarity of the leakage current. For example, a leakage current of 1 nA would cause the voltage on the loop capacitor (1000 pF for example) to charge or discharge by $dV/dt = I/C$ (1 V/s in this case). This, in turn, would cause the VCO to drift. So, if the loop is open for 1 ms and the K_V of the VCO is 50 MHz/Volt, the frequency drift caused by 1-nA leakage into a 1000-pF loop capacitor would be 50 kHz. In fact, the DECT bursts are generally shorter (0.5 ms), so the drift will be even less in practice for the loop capacitance and leakage current used in the example. However, it does serve to illustrate the importance of charge-pump leakage in this type of application.

Receiver Sensitivity

Receiver sensitivity specifies the ability of the receiver to respond to a weak signal. Digital receivers use maximum bit-error rate (BER) at a certain rf level to specify performance. In general, device gains, noise figures, image noise, and local-oscillator (LO) wideband noise all combine to produce an equivalent noise figure. This is then used to calculate the overall receiver sensitivity.

Wideband noise in the LO can elevate the IF noise level and thus degrade the overall noise factor. For example, wideband phase noise at $F_{LO} + F_{IF}$ will produce noise products at F_{IF} . This directly impacts the receiver sensitivity. This wideband phase noise is primarily dependent on the VCO phase noise.

Close-in phase noise in the LO will also impact sensitivity. Obviously, any noise close to F_{LO} will produce noise products close to F_{IF} and impact sensitivity directly.

Receiver Selectivity

Receiver selectivity specifies the tendency of a receiver to respond to channels adjacent to the desired reception channel. *Adjacent-channel interference* (ACI), a commonly used term in wireless systems, is also used to describe this phenomenon. When considering the LO section, the reference spurs are of particular importance with regard to selectivity. Figure 13 is an attempt to illustrate how a spurious signal at the LO, having the same spacing as the channel-spacing frequency, can translate energy from an adjacent radio channel directly onto the IF. This is of particular concern if the desired received signal is distant and weak and the unwanted adjacent channel is nearby and strong, which can often be the case. So, the lower the reference spurs in the PLL, the better it will be for system selectivity.

Conclusion

In Part 2 of this series we have discussed some of the critical specifications associated with PLL synthesizers, described measurement techniques, and shown examples of results. In addition, there has been a brief discussion of the system implications of phase noise, reference spurs and leakage current.

In the final part of this series, we will examine the building blocks that go to make up a PLL synthesizer. In addition, there will be a comparison between integer-N and fractional-N architectures for PLL.

Acknowledgment

The authors would like to acknowledge Brendan Daly of the Analog Devices GP RF Applications Group in Limerick for the plots of phase noise and reference spurs.

References

1. Mini-Circuits Corporation, *VCO Designers' Handbook*, 1996.
2. L.W. Couch, *Digital and Analog Communications Systems*, Macmillan Publishing Company, New York, 1990.
3. P. Vizmuller, *RF Design Guide*, Artech House, 1995.
4. R.L. Best, *Phase Locked Loops: Design, Simulation and Applications*, 3rd edition, McGraw-Hill, 1997.
5. D.E. Fague, "Open Loop Modulation of VCOs for Cordless Telecommunications," *RF Design*, July 1994.
6. D. B. Leeson, "A Simplified Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, Volume 42, February 1965, pp. 329-330. ▀




Figure 13. Adjacent Channel Interference.

Phase Locked Loops for High-Frequency Receivers and Transmitters—Part 3

Mike Curtin and Paul O'Brien

The first part of this series introduced phase-locked loops (PLLs), described basic architectures and principles of operation. It also included an example of where a PLL is used in communications systems. In the second part of the series, critical performance specifications, like phase noise, reference spurs, and output leakage, were examined in detail, and their effects on system performance were considered. In this, the last part of the series, we will deal with some of the main building blocks that go to make up the PLL synthesizer. We will also compare integer-N and fractional-N architectures. The series will end with a summary of VCOs currently available on the market and a listing of the Analog Devices family of synthesizers.

PLL Synthesizer Basic Building Blocks

A PLL synthesizer can be considered in terms of several basic building blocks. Already touched upon, they will now be dealt with in greater detail:

- Phase-Frequency Detector (PFD)
- Reference Counter (R)
- Feedback Counter (N)

The Phase-Frequency Detector (PFD)

The heart of a synthesizer is the phase detector—or phase-frequency detector. This is where the reference frequency signal is compared with the signal fed back from the VCO output, and the resulting error signal is used to drive the loop filter and VCO. In a digital PLL (DPLL) the phase detector or phase-frequency detector is a logical element. The three most common implementations are :

- Exclusive-or (EXOR) Gate
- J-K Flip-Flop
- Digital Phase-Frequency Detector

Here we will consider only the PFD, the element used in the ADF4110 and ADF4210 synthesizer families, because—unlike the EXOR gate and the J-K flip flop—its output is a function of both the frequency difference and the phase difference between the two inputs when it is in the unlocked state.

Figure 1 shows one implementation of a PFD, basically consisting of two D-type flip flops. One Q output enables a positive current source; and the other Q output enables a negative current source. Assuming that, in this design, the D-type flip flop is positive-edge triggered, the states are these (Q1, Q2):

- 11—both outputs high, is disabled by the AND gate (U3) back to the CLR pins on the flip flops.
- 00—both P1 and N1 are turned off and the output, OUT, is essentially in a high impedance state.
- 10—P1 is turned on, N1 is turned off, and the output is at V+.
- 01—P1 is turned off, N1 is turned on, and the output is at V-.




Figure 1. Typical PFD using D-type flip flops.

Consider now how the circuit behaves if the system is out of lock and the frequency at +IN is much higher than the frequency at -IN, as exemplified in Figure 2.




Figure 2. PFD waveforms, out of frequency and phase lock.

Since the frequency at +IN is much higher than that at -IN, the output spends most of its time in the high state. The first rising edge on +IN sends the output high and this is maintained until the first rising edge occurs on -IN. In a practical system this means that the output, and thus the input to the VCO, is driven higher, resulting in an increase in frequency at -IN. This is exactly what is desired.

If the frequency on +IN were much lower than on -IN, the opposite effect would occur. The output at OUT would spend most of its time in the low condition. This would have the effect of driving the VCO in the negative direction and again bring the frequency at -IN much closer to that at +IN, to approach the locked condition. Figure 3 shows the waveforms when the inputs are frequency-locked and close to phase-lock.




Figure 3. PFD waveforms, in frequency lock but out of phase lock.

Since +IN is leading -IN, the output is a series of positive current pulses. These pulses will tend to drive the VCO so that the -IN signal become phase-aligned with that on +IN.

When this occurs, if there were no delay element between U3 and the CLR inputs of U1 and U2, it would be possible for the output to be in high-impedance mode, producing neither positive nor

negative current pulses. This would not be a good situation. The VCO would drift until a significant phase error developed and started producing either positive or negative current pulses once again. Over a relatively long period of time, the effect of this cycling would be for the output of the charge pump to be modulated by a signal that is a subharmonic of the PFD input reference frequency. Since this could be a low frequency signal, it would not be attenuated by the loop filter and would result in very significant spurs in the VCO output spectrum, a phenomenon known as the *backlash* effect. The delay element between the output of U3 and the CLR inputs of U1 and U2 ensures that it does not happen. With the delay element, even when the +IN and -IN are perfectly phase-aligned, there will still be a current pulse generated at the charge pump output. The duration of this delay is equal to the delay inserted at the output of U3 and is known as the *anti-backlash pulse width*.

The Reference Counter

In the classical Integer-N synthesizer, the resolution of the output frequency is determined by the reference frequency applied to the phase detector. So, for example, if 200-kHz spacing is required (as in GSM phones), then the reference frequency must be 200 kHz. However, getting a stable 200-kHz frequency source is not easy. A sensible approach is to take a good crystal-based high frequency source and divide it down. For example, the desired frequency spacing could be achieved by starting with a 10-MHz frequency reference and dividing it down by 50. This approach is shown in the diagram in Figure 4.




Figure 4. Using a reference counter in a PLL synthesizer.

The Feedback Counter, N

The N counter, also known as the N divider, is the programmable element that sets the relationship between the input and output frequencies in the PLL. The complexity of the N counter has grown over the years. In addition to a straightforward N counter, it has evolved to include a prescaler, which can have a dual modulus.

This structure has grown as a solution to the problems inherent in using the basic divide-by-N structure to feed back to the phase detector when very high-frequency outputs are required. For example, let's assume that a 900-MHz output is required with 10-kHz spacing. A 10-MHz reference frequency might be used, with the R-Divider set at 1000. Then, the N-value in the feedback would need to be of the order of 90,000. This would mean at least a 17-bit counter capable of dealing with an input frequency of 900 MHz.

To handle this range, it makes sense to precede the programmable counter with a fixed counter element to bring the very high input frequency down to a range at which standard CMOS will operate. This counter, called a *prescaler*, is shown in Figure 5.

However, using a standard prescaler introduces other complications. The system resolution is now degraded ($F_1 \times P$). This issue can be addressed by using a dual-modulus prescaler (Figure 6). It has the advantages of the standard prescaler but without any loss in system resolution. A dual-modulus prescaler is a counter whose




Figure 5. Basic prescaler.

division ratio can be switched from one value to another by an external control signal. By using the dual-modulus prescaler with an A and B counter one can still maintain output resolution of F_1 . However, the following conditions must be met:

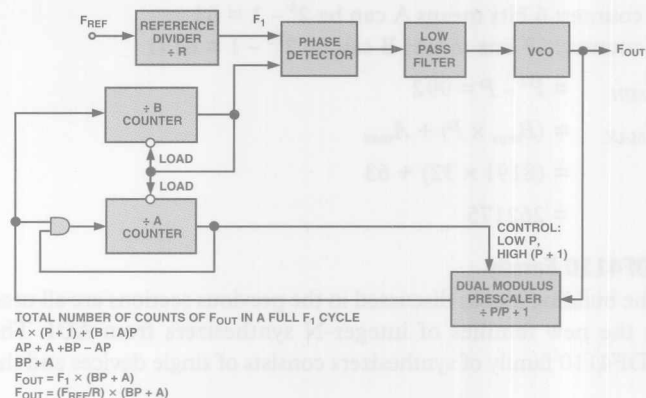


Figure 6. Dual-modulus prescaler.

1. The output signals of both counters are High if the counters have not timed out.
2. When the B counter times out, its output goes Low, and it immediately loads both counters to their preset values.
3. The value loaded to the B counter must always be greater than that loaded to the A counter.

Assume that the B counter has just timed out and both counters have been reloaded with the values A and B. Let's find the number of VCO cycles necessary to get to the same state again.

As long as the A counter has not timed out, the prescaler is dividing down by $P + 1$. So, both the A and B counters will count down by 1 every time the prescaler counts ($P + 1$) VCO cycles. This means the A counter will time out after $((P + 1) \times A)$ VCO cycles. At this point the prescaler is switched to divide-by-P. It is also possible to say that at this time the B counter still has $(B - A)$ cycles to go before it times out. How long will it take to do this: $((B - A) \times P)$. The system is now back to the initial condition where we started.

The total number of VCO cycles needed for this to happen is :

$$N = (A \times (P + 1)) + ((B - A) \times P)$$

$$= AP + A + BP - AP$$

$$= A + BP$$

When using a dual-modulus prescaler, it is important to consider the lowest and highest values of N. What we really want here is the range over which it is possible to change N in discrete integer steps. Consider the expression $N = A + BP$. To ensure a continuous integer spacing for N, A must be in the range 0 to $(P - 1)$. Then, every time B is incremented there is enough resolution to fill in all the integer values between BP and $(B + 1)P$. As was already noted for the dual-modulus prescaler, B must be greater than or equal to A for the dual modulus prescaler to work. From these we can

say that the smallest division ratio possible while being able to increment in discrete integer steps is:

$$\begin{aligned} N_{MIN} &= (B_{min} \times P) + A_{min} \\ &= ((P-1) \times P) + 0 \\ &= P^2 - P \end{aligned}$$

The highest value of N is given by

$$N_{MAX} = (B_{max} \times P) + A_{max}$$

In this case A_{max} and B_{max} are simply determined by the size of the A and B counters.

Now for a practical example with the ADF4111.

Let's assume that the prescaler is programmed to 32/33.

A counter: 6 bits means A can be $2^6 - 1 = 63$

B counter: 13 bits means B can be $2^{13} - 1 = 8191$

$$N_{MIN} = P^2 - P = 992$$

$$\begin{aligned} N_{MAX} &= (B_{max} \times P) + A_{max} \\ &= (8191 \times 32) + 63 \\ &= 262175 \end{aligned}$$

ADF4110 Family

The building blocks discussed in the previous sections are all used in the new families of integer-N synthesizers from ADI. The ADF4110 family of synthesizers consists of single devices and the

ADF4210 family consists of dual versions. The block diagram for the ADF4110 is shown below. It contains the reference counter, the dual-modulus prescaler, the N counter and the PFD blocks described above.

Fractional-N Synthesizers*

Many of the emerging wireless communication systems have a need for faster switching and lower phase noise in the local oscillator (LO). Integer N synthesizers require a reference frequency that is equal to the channel spacing. This can be quite low and thus necessitates a high N. This high N produces a phase noise that is proportionally high. The low reference frequency limits the PLL lock time. Fractional-N synthesis is a means of achieving both low phase noise and fast lock time in PLLs.

The technique was originally developed in the early 1970s. This early work was done mainly by Hewlett Packard and Racal. The technique originally went by the name of "digiphase" but it later became popularly named fractional-N.

In the standard synthesizer, it is possible to divide the RF signal by an integer only. This necessitates the use of a relatively low reference frequency (determined by the system channel spacing) and results in a high value of N in the feedback. Both of these facts have a major influence on the system settling time and the system phase noise. The low reference frequency means a long settling time, and the high value of N means larger phase noise.

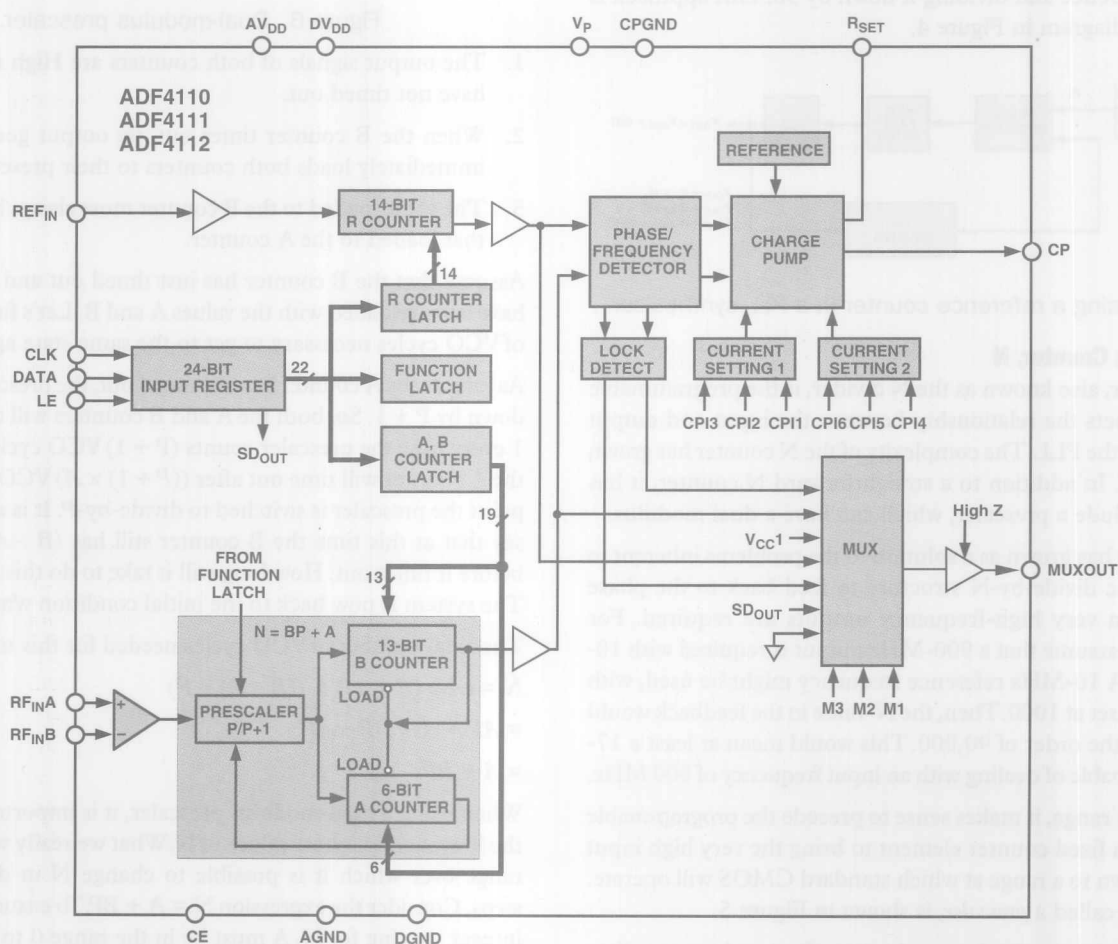


Figure 7. Block diagram for the ADF4110 family.

*The authors are indebted to The McGraw-Hill Companies for permission to use copyrighted material from Reference 4 in this section.

If division by a fraction could occur in the feedback, it would be possible to use a higher reference frequency and still achieve the channel spacing. This lower fractional number would also mean lower phase noise.

In fact it is possible to implement division by a fraction over a long period of time by alternately dividing by two integers (divide by 2.5 can be achieved by dividing successively by 2 and 3).

So, how does one divide by X or $(X + 1)$ (assuming that the fractional number is between these two values)? Well, the fractional part of the number can be allowed to accumulate at the reference frequency rate.

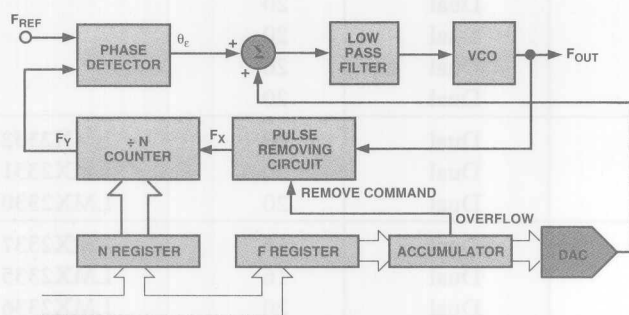


Figure 8. Fractional-N synthesizer.

Then every time the accumulator overflows, this signal can be used to change the N divide ratio. This is done in Figure 8 by removing one pulse being fed to the N counter. This effectively increases the divide ratio by one every time the accumulator overflows. Also, the bigger the number in the F register, the more often the accumulator overflows and the more often division by the larger number occurs. This is exactly what is desired from the circuit. There are some added complications however. The signal being fed to the phase detector from the divide-by- N circuit is not uniform in real time. Instead it is being modulated at a rate determined by the reference frequency and the programmed fraction. This in turn modulates the phase detector output and goes to the VCO input. The end result is a high spurious content at the output of the VCO. Major efforts are currently under way to minimize these spurs. One method uses the DAC shown in Figure 8.

Up to now, monolithic Fractional- N synthesizers have failed to live up to expectations but the eventual benefits that may be realized mean that development is continuing at a rapid pace.

Summary of VCO Manufacturers

With the explosive growth in wireless communications, the demand for products like synthesizers and VCOs has increased dramatically over the past five years. Interestingly, until now, the markets have been served by two distinct sets of manufacturers. Below is listed

a selection of players in the VCO field. This list is not meant to be all-inclusive, but rather gives the reader a feel for some of the main players.

VCOs

Murata	Murata has both 3-V and 5-V devices available. The VCOs are mainly narrowband for wireless handsets and base stations. Frequencies are determined by the wireless frequency standards.
Vari-L	Vari-L addresses the same market as Murata. 3-V and 5-V devices are available.
Alps	Alps makes VCOs for wireless handsets and base stations.
Mini-Circuits	Mini-Circuits offers both narrowband and wideband VCOs.
Z-Comm	Z-Communications has both wideband and narrowband VCOs. The wideband VCOs typically have an octave tuning range (1 GHz to 2 GHz, for example) and operate from a supply voltage of up to 20 V. They offer surface-mount packaging.
Micronetics	Micronetics offers both narrowband and wideband VCOs. Their strength lies more in the wideband products where they can go from an octave range at anything up to 1200 MHz. Above these output frequencies, the range is somewhat reduced.

The Analog Devices Synthesizer Family

The table on the next page lists current and future members of the ADF4xxx synthesizer family. It includes single and dual, and integer- N and fractional- N devices.

Acknowledgements

The ADF4xxx Family of synthesizers is designed at the Analog Devices facility in Limerick, Ireland. The product line team includes: Mike Tuthill, Leo McHugh, Bill Hunt, Mike Keaveney, Brendan Daly, Paul O'Brien, Paul Mallon, Ian Collins, Sinead O'Keefe, Liam McCann, Patrick Walsh, Cristoir O'Reilly, Paul Laven, Samuel Landete, Niall Kearney, and Mike Curtin. The group would like to acknowledge the valuable contributions of Jon Strange and Ashish Shah at Analog Devices, Kent (U.K.), and of Fred Weiss at Analog Devices Northwest Labs (Beaverton, OR).

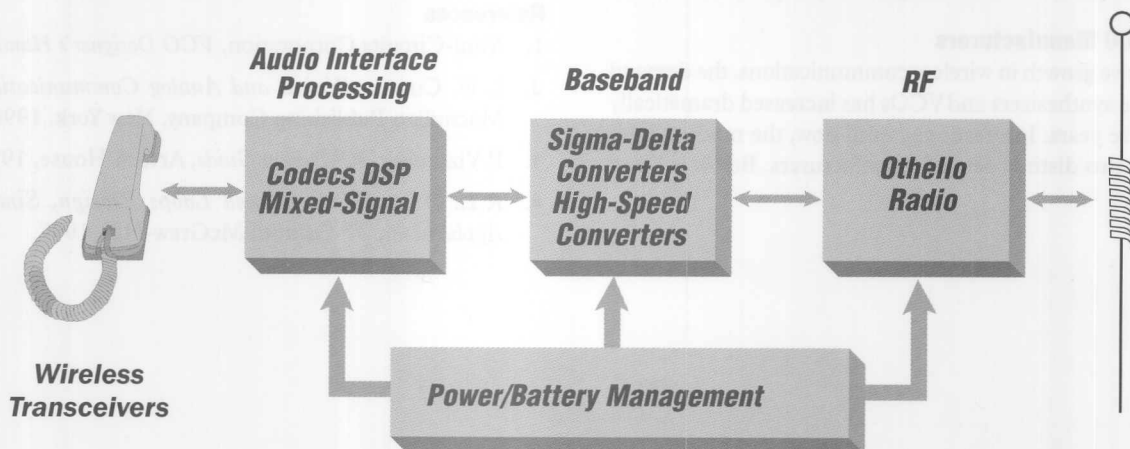
References

1. Mini-Circuits Corporation, *VCO Designer's Handbook*, 1996.
2. L.W. Couch, *Digital and Analog Communications Systems*, Macmillan Publishing Company, New York, 1990.
3. P. Vizmuller, *RF Design Guide*, Artech House, 1995.
4. R.L. Best, *Phase Locked Loops: Design, Simulation and Applications*, 3rd Edition, McGraw-Hill, 1997. ▀

ADF4xxx Frequency Synthesizer Family

Device	Integer-N Frequency Range	Fractional-N Frequency Range	Single/Dual	Pin Count	Second Source
ADF4110	≤550 MHz		Single	16	
ADF4111	≤1.2 GHz		Single	16	
ADF4112	≤3.0 GHz		Single	16	
ADF4113	≤3.8 GHz		Single	16	
ADF4116	≤550 MHz		Single	16	LMX2306
ADF4117	≤1.2 GHz		Single	16	LMX2316
ADF4118	≤3.0 GHz		Single	16	LMX2326
ADF4210	≤510 MHz/≤1.2 GHz		Dual	20	
ADF4211	≤510 MHz/≤2.0 GHz		Dual	20	
ADF4212	≤510 MHz/≤3.0 GHz		Dual	20	
ADF4213	≤1.0 GHz/≤2.5 GHz		Dual	20	
ADF4216	≤510 MHz/≤1.2 GHz		Dual	20	LMX2332L
ADF4217	≤510 MHz/≤2.0 GHz		Dual	20	LMX2331L
ADF4218	≤510 MHz/≤2.5 GHz		Dual	20	LMX2330L
ADF4206	≤500 MHz/≤500 MHz		Dual	16	LMX2337
ADF4207	≤1.1 GHz/≤1.1 GHz		Dual	16	LMX2335
ADF4208	≤1.1 GHz/≤2.0 GHz		Dual	20	LMX2336
ADF4150		≤550 MHz	Single	16	
ADF4151		≤1.2 GHz	Single	16	
ADF4152		≤3.0 GHz	Single	16	
ADF4156		≤550 MHz	Single	20	
ADF4157		≤1.2 GHz	Single	20	
ADF4158		≤3.0 GHz	Single	20	
ADF4250	≤550 MHz	≤1.2 GHz	Dual	20	
ADF4251	≤550 MHz	≤2.0 GHz	Dual	20	
ADF4252	≤1.0 GHz	≤3.0 GHz	Dual	20	
ADF4256		≤550 MHz/≤1.2 GHz	Dual	20	
ADF4257		≤550 MHz/≤2.0 GHz	Dual	20	
ADF4258		≤1.0 GHz/≤3.0 GHz	Dual	20	

All brand or product names mentioned are trademarks or registered trademarks of their respective holders.



Dual Axis, Low g, Fully Integrated Accelerometers

By Harvey Weinberg

The ADXL202 is the newest low- g ($\pm 2 g$), dual-axis, surface-micromachined accelerometer from Analog Devices. Building on experience gained in manufacturing millions of iMEMS[®] accelerometers in the past six years, the ADXL202 is the world's first commercial dual axis, surface micromachined accelerometer to combine low- g sensing with lowest power, lowest noise, and digital outputs—all on a single silicon chip.

Surface micromachining, first commercialized with the ADXL50, allows for integration of the acceleration sensor with all signal conditioning electronics—tight integration of the sensor and its signal conditioning is what has made this impressive performance possible.

Lower cost was a major driver in the ADXL202 design effort. Integrating two axes resulted in a significant cost reduction per axis. In addition, while the ADXL50, ADXL150, and ADXL250 can be thought of as “acceleration to volts” transducers, the ADXL202 adds a Pulse Width Modulated (PWM) digital output capability as well. Since most accelerometers will interface with a microcontroller, a PWM output obviates the need for an A to D converter, further driving down the user's total system cost.

Sensor Structure

As with all of our accelerometer products, the sensor element is a differential capacitor whose output is proportional to acceleration (basic sensor information can be seen in *Analog Dialogue* 27-2, 1993, and *Analog Dialogue* 30-4, 1996). Since device performance is so dependent on sensor design, a brief explanation of some of the key factors in beam design is appropriate.

The beam is made up of many interdigitated fingers. Each set of fingers can be visualized as shown in Figure 1. The differential capacitance of each finger is proportional to the overlapping area between the fixed outer plates and the moving finger, and the displacement of the moving finger. Clearly these are very small capacitors, and in order to reduce noise and increase resolution we need as large a differential capacitance as practical.




Figure 1. Beam Dimensions for a Single Finger.

The capacitor area is limited by the 2-micron height—fixed by process technology, while the (125-micron) overlap is adjustable to some extent. However longer fingers are not desirable for several reasons. Longer fingers are harder to manufacture and increased beam size translates to more expensive parts.

The movement of the beam is controlled by the polysilicon springs holding the beam. These springs and the beam's mass obey the same laws of physics we learned in high school. The force (F) on a mass (m) subject to acceleration (a), according to Newton's Second Law, is $F = m a$, and the deflection (x) of a restraining spring (obeying Hooke's Law) is proportional to the applied force, $F = k x$, and:

$$x = (m/k) a$$

The only two parameters under our control are the spring stiffness, or spring constant, (k), and mass (m). Reducing the spring constant seems like an easy way to improve beam sensitivity. But as usual, nothing comes for free. The resonant frequency of the beam is proportional to the spring constant, and the accelerometer must operate at frequencies below the resonant frequency. In addition, higher spring constants make for more rugged beams (higher shock survivability). So if we would like to keep the spring constant as high as possible the only parameter left to change is mass.

Adding mass normally implies a larger sensor area, resulting in more expensive parts, since the only way to add mass it to make the beam larger. In the ADXL202 a novel beam structure was invented, as shown in Figure 2. Rather than using two discrete beams placed orthogonally as in the ADXL250 (*Analog Dialogue* 30-4, 1996, page 5, Figure 5), the fingers that constitute the X and Y axis variable capacitors are integrated along the sides of a single square beam. This results in a reduction of the overall sensor area, yet the larger common beam mass enhances the resolution of the ADXL202. A spring suspension system, shown in Figure 3, situated in the corners of the beam, was designed to minimize cross-axis sensitivity (i.e., with acceleration along one axis, any tendency toward motion or outputs in the orthogonal direction is suppressed).




Figure 2. ADXL202 Beam Structure.




Figure 3. Detail of One Corner of the ADXL202 Beam showing springs.

Circuit Architecture

The circuit architecture (shown in Figure 4) of the ADXL202 is similar to the ADXL250 up to the demodulator. The fixed outer plates are driven with square waves that are 180 degrees out of phase. When the movable fingers (and hence the beam) are centered between the fixed outer plates, both sides of the differential capacitor have equal capacitance and the ac voltage on the beam is zero. However, if the beam is displaced due to an applied acceleration, the differential capacitance will be unbalanced and an ac voltage of amplitude proportional to the displacement of the beam will result.

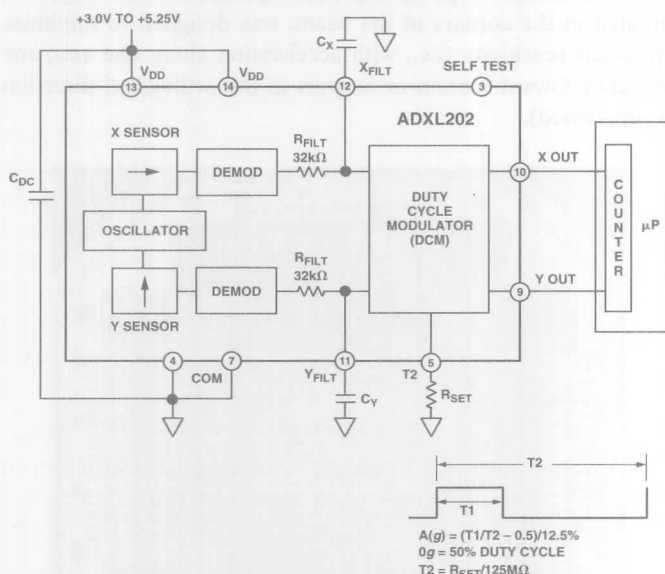


Figure 4. ADXL202 Circuit Architecture.

This ac voltage is amplified and then demodulated by a synchronous demodulator. The output of the demodulator drives the duty cycle modulator through a 32-kΩ resistor. Here a pin is available on each channel to allow the user to set the bandwidth by adding two external capacitors (one per channel) creating a simple first-order RC low pass filter. The low pass filtered signal is converted to a PWM signal by the duty cycle modulator. The period of the PWM output may be set from 0.5 to 10 ms, using a single resistor.

Performance and Applications

Just as it was impossible 25 years ago to predict where low-cost lasers would turn up, today it's difficult to imagine the large number of applications where low-cost accelerometers will be used. The ADXL202 breaks so many performance-vs.-cost barriers that most of its really successful applications are not classically "accelerometer" (literally "acceleration-measuring") applications. They are now being used in car alarms, machine health monitoring, joysticks, game pads, and other computer input devices.

As mentioned in the introduction, the ADXL202 is the lowest noise dual axis surface micromachined accelerometer in production today. With its typical noise density of 500 mg/√Hz, it is possible to resolve inclinations of better than $\pm 1^\circ$ of tilt at bandwidths of up to 50 Hz. The high-resolution (approximately 14-bit) duty-cycle modulator allows users to take advantage of the ADXL202's capabilities in low cost systems. These capabilities have opened the door to several other non-traditional applications, such as car alarms (where they are used to sense jacking-up or towing) and automatic machine leveling.

Support Tools

Extensive support tools are available for designers. The hardware tools available are a simple evaluation PCB with an ADXL202 in its minimum circuit configuration (part number ADXL202EB), and the ADXL202EB-232, a complete 2-axis data acquisition system that interfaces to a PC. Primarily targeted at designers who need to fully understand how acceleration measurement will enable their application or product, the ADXL202EB-232 includes software for viewing accelerometer signals and data logging.

The ADXL202 Interactive Designer is an Excel-based spreadsheet model that takes the user through the design process of selecting ADXL202 components and defining software parameters for microcontroller interface. The spreadsheet outputs component values and information about the resolution, bandwidth and acquisition rate of your design. It is available free at the Web site listed below.

Also available at that site are several hardware and software reference designs, highlighting different interface techniques and a variety of application notes. Each reference design includes flow charts and source code for various popular microcontrollers, as well as a description of where each data acquisition method is appropriate. These documents can be found at http://www.analog.com/iMEMS/products/ADXL202_top.html. □

Ask the Applications Engineer—29

By Harvey Weinberg

ACCELEROMETERS—FANTASY & REALITY

Q. *As an Applications Engineer for the Micromachined Products Division supporting the compact, low-cost, gravity-sensitive Analog Devices line of integrated accelerometers¹ you must get to hear lots of creative ideas from prospective customers about how to employ accelerometers in useful ways.*

A. Yes, but sometimes what they are suggesting violates physical laws! We've come to rate them in various categories on an informal "reality" scale:

- **Real** – A real application that actually works today and is currently in production.
- **Possible** – An application that is technically feasible, but not (to our knowledge) in production.
- **Fantasy** – An application that could be possible if we had much better technology.
- **Dream Land** – An application that would be nice, but any practical implementation we can think of would violate some physical laws.

Would you like to see some of these ideas (the ones we're free to mention), hear about their potential benefits, think about them, and guess which category they belong in?

Q. *Sure, go ahead.*

A. O.K. You can find (and link to) the answers, with commentary, below.

Washing Machine Load Balancing

Unbalanced loads during the high-speed spin-cycle cause washing machines to shake, and in some cases, to even "walk" across the floor if unrestrained. An accelerometer senses the accelerations present during the unbalanced spin. If such an imbalance is present, the washing machine redistributes the load by jogging the drum back and forth a few times until the load is balanced. *Real or fantasy?* See answer A.

Machine Health Monitors

Many industries change or overhaul mechanical equipment using a calendar-based preventive maintenance schedule. This is especially true in applications where one cannot afford or tolerate unscheduled down-time. So machinery with plenty of service life left is often prematurely rebuilt at a cost of millions of dollars across many industries. By embedding accelerometers in bearings, or other rotating equipment, service life can be extended without additional risk of sudden failure.

The accelerometer senses the vibration of bearings or other rotating equipment to determine their condition. *Real or fantasy?* See answer B.

Automatic Leveling

Accelerometers measure the absolute inclination of an object (e.g. a large machine, a mobile home, etc.). The tilt information is used by a microcontroller to automatically level it. *Real or fantasy?* See answer C.

Scroll Control for PDAs and Handyphones

The accelerometer allows the microcontroller to recognize gestures allowing the user one handed control of mobile devices. *Real or fantasy?* See answer D.

Car Alarm

Here the accelerometer senses if a car is being jacked up or being picked up by a tow truck, and sets off the alarm. *Real or fantasy?* See answer E.

Ski Bindings

The accelerometer measures the total shock energy and signature to determine if the binding should release. *Real or fantasy?* See answer F.

Personal Navigation

In this application, position is determined by dead reckoning (double integration of acceleration over time to determine actual position). *Real or fantasy?* See answer G.

Subwoofer Servo Control

An accelerometer is mounted on the cone of the subwoofer to provide positional feedback used to servo out distortion. *Real or fantasy?* See answer H.

Neuromuscular Stimulator

This application helps people, who have lost control of their lower leg muscles, to walk—by stimulating muscles at the appropriate time. *Real or fantasy?* See answer I.

Car-Noise Cancellation

The accelerometer senses low-frequency vibration in the passenger compartment, and the noise-cancellation system nulls it out, using the speakers in the car stereo system. *Real or fantasy?* See answer J.

ANSWERS

A. *Washing-machine load balancing: Real.* This application is currently in production. With better load balance, faster spin rates can be used to wring more water out of clothing, making the drying process more efficient. Fewer mechanical components are required for damping the drum motion, making the overall system lighter, and less expensive. In addition, transmission and bearing service life may be extended because of lower peak loads present on the motor. See http://www.analog.com/industry/iMEMS/markets/industrial/washing_machine.html for more information about washing machines. *Go to next question.*

B. *Machine health monitoring: Possible.* This application has been demonstrated but is not yet in production.

Using the vibrational "signature" of bearings to determine their condition is a well proven and industry-accepted method of equipment maintenance. However, the cost of accelerometers and the associated signal conditioning equipment has been traditionally too high. The ADXL105² offers a complete vibration measurement and signal conditioning solution on a chip at very low cost. See <http://www.analog.com/industry/iMEMS/markets/industrial/machine.html> for more information about machine health. *Go to next question.*

¹<http://www.analog.com/industry/iMEMS>.

²<http://products.analog.com/products/info.asp?product=ADXL105>.

C. Automatic leveling: Real to Fantasy (depending on the application). There are some applications where this is practical and currently in production. Others are too demanding for current products.

Self-leveling is a very demanding application, as absolute precision is required. Surface micromachined accelerometers have impressive resolution, but absolute tilt measurement with high accuracy (to better than 1%) requires temperature stability and hysteresis performance that today's surface-micromachined accelerometers cannot achieve. Applications needing absolute accuracy to within $\pm 3^\circ$ or more are currently possible and a few such applications are in production. *Go to next question.*

D. Scroll control: Real. This application is currently (or almost, depending on press time) in production.

A PDA (like the 3com Palm Pilot) is incredibly handy, but almost impossible to use one-handed. Like when you're driving, or on the phone. Adding an accelerometer lets the PDA accept gesture inputs, like tilting up or down, to control the cursor or page up/down control. See <http://www.analog.com/industry/iMEMS/markets/consumer/peripherals.html> if you are interested in game controllers or are a Palm Pilot user and want to see how to add a tilt function to your PDA. *Go to next question.*

E. Car alarm: Real. This application is currently in production in OEM and after-market automotive anti-theft systems.

One of the most popular methods of auto theft is where the car is stolen by simply towing it away. Conventional car alarms do not protect against this. Shock sensors cannot measure changes in inclination, and ignition-disabling systems are ineffectual. Here is an application where the high-resolution capabilities of the ADXL202 are used to advantage. The accelerometer measures if the car's inclination is changing by more than 0.5° per minute. If so, the alarm is sounded, hopefully scaring off the would-be thief. Absolute stability is not required here (unlike automatic leveling systems) as temperature does not change significantly in a minute or less. See http://www.analog.com/industry/iMEMS/markets/automotive/car_alarms.html if you're interested in car alarms. *Go to next question.*

F. Ski bindings: Fantasy. This is a practical accelerometer application, but current battery technology (particularly low temperature performance) makes it impractical.

All mechanical ski bindings are highly evolved, but limited in their performance. Measuring the actual shock experienced by the skier would be a much more accurate way to determine if a binding should release. Intelligent systems could be developed that could take each individual's capability and physiology into account. Smaller and lighter batteries that perform well at low temperatures will, eventually, enable this application. *Go to next question.*

G. Personal navigation: Dream Land. Long term integration results in the accumulation of error due to small dc errors in the accelerometer, integrator input circuitry, wiring thermocouples, etc. Double integration compounds the errors (t^2). Without some way of "resetting" the actual position from time to time, huge errors result. This is analogous to building

an op-amp integrator by simply putting a capacitor across it. Even if the accelerometer's accuracy is improved by ten or one hundred times better than currently available, huge errors would still eventually result. It would just take longer to happen.

Accelerometers can be used in conjunction with a GPS system when the GPS signals are briefly unavailable. Integration over a short time (a minute or so) can give satisfactory results. See http://www.analog.com/industry/iMEMS/markets/consumer/car_nav.html for more information about navigation. *Go to next question.*

H. Subwoofer servo: Real. Several active subwoofers with servo control are on the market today.

Servo control of subwoofers has several advantages. Harmonic distortion, as well as power compression, can be greatly reduced. In addition servo control can also electronically lower the Q of the speaker/enclosure system, enabling the use of smaller enclosures³. The ADXL190⁴ is small and light; its mass, added to that of the loudspeaker cone, does not change the overall acoustic characteristics significantly. See <http://www.analog.com/industry/iMEMS/markets/consumer/subwoofers.html> for more information about active subwoofer applications, and <http://www.analog.com/industry/iMEMS/markets/consumer/subwoofers/Subwoof.html#cir> for circuits. *Go to next question.*

I. Neuromuscular stimulator: Real. This application is very near (if not already in) production.

When walking, the forefoot is normally raised when moving the leg forward, then lowered when pushing the leg backward. The accelerometer is worn somewhere on the lower leg or foot, where it senses the position of the leg. The appropriate muscles are then electronically stimulated to flex the foot as required.

This is a classic example of how micromachined accelerometers have made a product feasible. Earlier models used a liquid tilt sensor or a moving ball bearing (acting as a switch) to determine the leg position. Liquid tilt sensors had problems because of sloshing of the liquid, so only slow walking was possible. Ball-bearing switches were easily confused when walking on hills. Using an accelerometer, the differential between leg back and leg forward is measured, so hills do not fool the system and there are no liquid slosh problems. The low power consumption of the accelerometer allows the system to work with a small lithium battery, making the overall package unobtrusive. *Go to next question.*

J. Car-noise cancellation: Dream Land. While the accelerometer has no trouble picking up the vibration in the passenger compartment, noise cancellation is highly phase-dependent. So while we may cancel the noise in one location (say around the head of the driver), it will probably be increased at other locations.

Conclusion

Because of their sensitivity, compactness, low cost, ruggedness, and ability to measure both static and dynamic acceleration forces, surface micromachined accelerometers have made numerous new applications possible. Many of them were not anticipated because they were not thought of as classic accelerometer applications. The imagination of designers now seems to be the limiting factor in the scope of potential applications—but sometimes designers can become too imaginative! While performance improvements continue to enable more applications, it's wise to try to stay away from "solutions" that violate laws of physics. ■

³See R. A. Greiner and T. M. Sims, Jr., Loudspeaker Distortion Reduction, *Journal of the Audio Engineering Society*, Vol. 32, No. 12.

⁴<http://products.analog.com/products/info.asp?product=ADXL190>.

ADXL105: A Lower Noise, Wider Bandwidth Accelerometer Rivals Performance of More Expensive Sensors

by James Doscher, Marketing Manager,
Micromachined Product Division

INTRODUCTION

Surface micromachining has been used to make reliable accelerometers in large quantities at very low cost, but with modest performance, for the airbag and consumer products industry (*Analog Dialogue* 27-2 (1993), *A-D* 30-4 (1996), *A-D* 33-1 (1999)). The accuracy and noise limitations for more-sophisticated applications are due to the tiny signals involved, given the small size of surface-micromachined structures.

However, recent advances in circuit architectures and beam structures used in integrated micro-electromechanical systems have resulted in order-of-magnitude improvements in resolution and accuracy. The recently introduced ADXL105 accelerometer has a $225\text{-}\mu\text{g}/\sqrt{\text{Hz}}$ noise floor, 10-kHz bandwidth, and an on-board temperature sensor, which can be used for calibrating against temperature effects. In comparison, the ADXL50, the first Analog Devices accelerometer, had a noise floor of $6500\text{ }\mu\text{g}/\sqrt{\text{Hz}}$. A $65\text{-}\mu\text{g}/\sqrt{\text{Hz}}$ version has also been demonstrated on ADI's new $3\text{-}\mu\text{m}$ polysilicon process. The ADXL105 is initially targeted at machine health and tactical guidance & control applications, as well as a number of new automotive applications.

The ADXL105, a near ideal kind of vibration sensor, eliminates significant problems with existing vibration sensors, such as piezoelectric and bulk capacitive sensors. Primary benefits derive from much lower cost, stable sensitivity as a function of frequency and temperature, ruggedness, and ease of use. Besides machine health and condition monitoring, it is particularly well suited for noise and vibration cancellation applications.

Structurally, the ADXL105 is the first of a new generation of precision surface-micromachined devices that attempts to push the performance envelope using a $2\text{-}\mu\text{m}$ -thick polysilicon process. The device is fabricated on a process that combines the surface-micromachined differential-capacitance sensor and the electronic signal conditioning circuitry on a single IC chip. The sensor itself is $0.5\text{ mm} \times 0.4\text{ mm}$ and $2\text{-}\mu\text{m}$ -thick, with $1\text{-}\mu\text{m}$ feature size and a mass of only 0.5 micrograms. The deflection due to a 1-g acceleration is 1 nanometer ($E\text{-}9\text{ m}$), or about 0.07% of the gap width, and the minimum resolvable deflection is about 210 femtometers ($E\text{-}15\text{ m}$), or 0.002 Angstrom units, at a noise level of $225\text{ }\mu\text{g}/\sqrt{\text{Hz}}$ and 1-Hz bandwidth.

With a source capacitance in the neighborhood of 150 femtofarads (i.e., 0.15 pF), and orders-of-magnitude smaller differential capacitance changes (100 attofarads ($E\text{-}18\text{ F}$) for a 1-g acceleration and 23 zeptofarads ($E\text{-}21\text{ F}$) resolution, the signal output of these

devices could be easily lost in the presence of parasitic capacitance or noise.

The ability to provide on-chip signal conditioning is what allows the minute signals to be read with minimal interference. A synchronous demodulation technique ("ac bridge") is used to increase noise immunity and improve signal to noise. A square-wave carrier drives the sensor, and the modulated signal produced by the change in differential capacitance is amplified and synchronously demodulated, minimizing noise and boosting the signal-to-noise ratio.

The limitations on size and mass of the structural beam place a weighty constraint on the performance of these devices. To make a precision accelerometer, it is helpful to have as much mass as possible and a large output signal from the sensor. But both parameters are limited by the size of the sensor element, which in turn is limited by the ability to control the curvature, or out-of-plane deviation of the structure due to internal polysilicon stress. Limits on size (hence mass) set a lower noise limit on the sensor, because the Brownian noise floor of the device, which is imposed by gas molecules bouncing off the beam, is inversely proportional to the mass of the sensor. Signal to noise is of course limited by the size of the signal, (related to spring stiffness, mass of the sensor, and source capacitance), as well as the noise floor set by the front-end circuitry.

Spring stiffness limits capacitance change, hence the magnitude of the electrical output. Although more-flexible springs allow greater travel of the mass under acceleration, with a corresponding increase in the delta capacitance read by the circuit, there is a direct trade-off between the flexibility of the springs and the robustness of the structure to shock and overload. If sensors are insufficiently stiff, their elements can touch, and the resulting stiction (static friction) produces a failure through irreversible physical latchup.

The design goal of the ADXL105 was to improve the robustness of the sensor, and at the same time achieve an approximately $5\times$ improvement in performance over existing products built in the process; this would improve the suitability of the devices for machine health and tactical inertial applications.

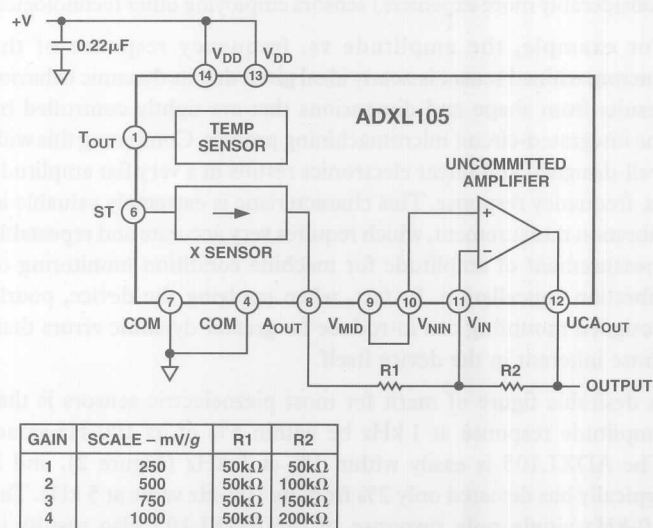


Figure 1. ADXL105 and typical application circuit.

To improve the robustness, and meet a target specification of 10-kHz bandwidth, a stiffer beam was selected. Supplanting a beam with a 12-kHz resonance, the new design has a 16-kHz resonant frequency, representing a 2.5× improvement in stiffness. In addition, the beam has a new suspension with increased Z-axis (beam to substrate) stiffness. While robustness was improved, the changes in fact increased the difficulty of meeting the design goal, as the stiffer beam reduced the signal coming from the sensor. This called for improvements in electrical design.

An especial effort was made to reduce or eliminate all sources of inaccuracy due to the electronics. Thus care was taken to use low drift amplifier designs and to track down error sources that were not important to airbag or consumer electronic designs, but would compromise an inertial or vibration sensor.

The front-end circuitry of the ADXL105 was designed using familiar precision design techniques to reduce the noise floor. Drive current and input FET size were increased to minimize noise and parasitic loading. Die size and current consumption were compromised somewhat in favor of performance. A special two-stage amplifier design was used that split up the gain segments to optimize the noise and drift components.

The increased performance produced by electrical and mechanical design improvements has answered two frequent criticisms of surface micromachined devices for precision measurement: that they have limited frequency response and that their noise floor is too high and hence resolution is too low. The table below compares the ADXL105 with other low-g ADI accelerometers.

Specifications	ADXL105	ADXL202	ADXL05	ADXL50
Range	$\pm 5 g$	$\pm 2 g$	$\pm 5 g$	$\pm 50 g$
Noise ($\mu g/\sqrt{Hz}$)	225	500	500	6500
Bandwidth (kHz)	10	5	5	6
Supply Current (mA)	2	0.6	10	10
Number of Axes	1	2	1	1
Output Type	Analog	Analog/ Digital	Analog	Analog

Beyond these specifications, there are other favorable characteristics of the surface micromachined approach that result in a sensor that rivals or exceeds the performance of existing (and considerably more expensive) sensors employing other technologies.

For example, the amplitude vs. frequency response of the micromachined sensor is nearly ideal given that its dynamic behavior results from shape and dimensions that are tightly controlled by the integrated-circuit micromachining process. Combining this with well-designed and linear electronics results in a very flat amplitude vs. frequency response. This characteristic is extremely valuable in vibration measurement, which requires very accurate and repeatable measurement of amplitude for machine condition monitoring or vibration cancellation. In fact, when applying the device, poorly designed mounting can introduce far greater dynamic errors than those inherent in the device itself.

A desirable figure of merit for most piezoelectric sensors is that amplitude response at 1 kHz be within 5% of its 100-Hz value. The ADXL105 is easily within 1% at 1 kHz (Figure 2), and it typically has deviated only 2% from its 100-Hz value at 5 kHz. The 10-kHz single pole response of the ADXL105 also results in predictable phase response and low phase lag at 5 kHz, an important characteristic for noise and vibration cancellation schemes.




Figure 2. Amplitude response deviation.

In this type of sensor, as mentioned earlier, the motion of the beam in response to the force of acceleration is detected by a capacitive technique. The measured capacitance change ratio is nearly invariant with temperature, resulting in a sensitivity (V/g) that changes only $\pm 1\%$ over the -40 to $+85^\circ\text{C}$ industrial temperature range (Figure 3).




Figure 3. Sensitivity error as a function of temperature (uncompensated).

Finally, the noise characteristics of the ADXL105 are very different from those of the popular piezoelectric sensors commonly used for vibration. The synchronous demodulation scheme used to decode the amplified capacitive sensor output results in a noise spectral density that is essentially independent of frequency. This contrasts with the noise from a piezoelectric sensor, which increased substantially as the frequency decreases (Figure 4).




Figure 4. Noise floor comparison.

Displacement is an important characteristic to measure for low-speed equipment, such as large air conditioning fans. Excessive displacement is a sign of bearing wear and of a potential for mechanical failure. The vibration acceleration signal is integrated twice to compute displacement. The ADXL105 can measure acceleration all the way down to dc; piezoelectric sensors cannot. They lose sensitivity in the neighborhood of 1 Hz due to a low-frequency zero. The dc ability and excellent low-frequency noise characteristics of the ADXL105 enable the accurate measurement of vibration displacement at low frequencies.

Applications Issues in Vibration Measurement

The ADXL105 is significantly different from other vibration sensors in one final area: packaging. Most vibration sensors are in a packaged format, with stud or screw mounts, and a pigtail for electrical connections. The packages are hardened for industrial use, stainless steel is often used. The ADXL105 comes housed in a standard ceramic surface-mount IC package. For industrial applications, it must usually be repackaged and potted into a form appropriate for the target equipment and the type of acceleration measurement required. Analog Devices works with a number of 3rd party suppliers, who can provide the product in a variety of form factors.

As noted earlier, the near-ideal characteristics of the sensor can be compromised by careless mounting. When potting the accelerometer, it is important to consider the resonant frequencies of the IC package leads and the PC board it is mounted to. For example the leads of the surface mount package have a resonant frequency of about 7 kHz. This resonance could degrade accuracy in high-frequency vibration measurement. To solve the problem, the package should be glued to the PC board (if a PC board is used) in order to make a tight mechanical coupling from the sensor to the board. PC-board resonances typically occur at even lower frequencies and can be damped by potting the assembly in epoxy or using a thicker PC board.

Special Features for Inertial Measurements

Although the ADXL105 is optimized for vibration measurement, its dc accuracy is also improved over that of earlier ADI acceleration sensors. Special attention was paid to the design of the electronics to eliminate dc drift and noise sources that were ignored, or were simply not important in the applications the earlier products were used for. Particular attention was paid to the final output stage to lower the 1/f noise corner that limits stability at dc and low frequencies.

The accelerometer includes an 8-mV/°C on-board temperature sensor that measures the actual temperature of the IC die. For inertial applications, the accelerometer is mathematically modeled over temperature to eliminate static errors. The on-board temperature sensor is used both for initial calibrations of the sensor and as an input to the calibration model. The sensor was designed to be repeatable and accurate to reduce residuals in the modeling process.

Future Work

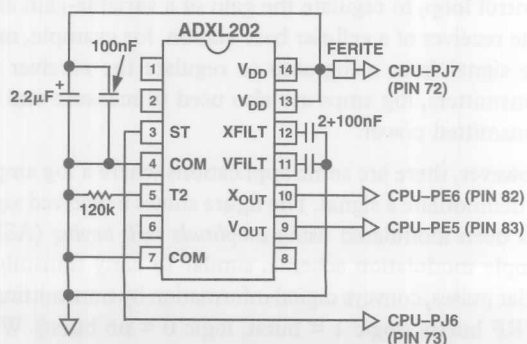
The ADXL105 has demonstrated that, despite inherent obstacles, it is possible to develop high-performance acceleration sensors on a 2- μm thick, surface-micromachine polysilicon process. An experimental version of the device, built on a 3- μm -thick polysilicon process, has demonstrated a $65\text{-}\mu\text{g}/\sqrt{\text{Hz}}$ noise floor, a $3\times$ improvement over the ADXL105. Concepts are now being developed to demonstrate a $25\text{-}\mu\text{g}/\sqrt{\text{Hz}}$ device, and to combine this increased resolution with wider ranges of $\pm 50\text{ g}$, $\pm 100\text{ g}$, and more. Other developments will focus on lower-stress packaging methods, which will enable better bias stability—an important requirement for tactical inertial applications.

Conclusions

The primary opportunity posed by the low cost and increasingly high performance of these devices is to enable new applications that weren't possible before. Although there is also a potential for replacing other means in existing applications, there is no reason to believe that designers who design or customers who purchase precision instruments will suddenly change to a new technology strictly on the basis of cost, as there are other important considerations in the precision instrumentation business.

However, there is a limitless opportunity to develop new or expanded uses of precision sensors. For example, the new technology may allow designers to instrument full-time monitoring of every machine, motor, pump or compressor that a manufacturer makes. New low cost, precision sensors with their fully signal conditioned output can reduce the cost of each point in a real-time monitoring system from \$1000 to under \$100. That kind of cost reduction makes feasible new uses for sensors that were not possible before. For example, a manufacturer of industrial equipment can now economically produce a "smart motor," which has high-efficiency drive electronics employing DSP for speed control, and accompany it by a low-cost machine-health conditioning module to increase energy efficiency, improve up time, reduce maintenance cost, and thus lower the end customer's total cost of ownership. ▢

ADDING A TILT SENSOR TO THE PALM PILOT



The AD202 can be added to a PalmPilot to measure acceleration and gravity. Read about "Adding a tilt sensor to the PalmPilot" at www.ibr.cs.tu-bs.de/~harbaum/pilot/adxl202.html. The modified PalmPilot can be used to play Mulg, a simple marble game. All you need to do is to guide the black marble to a little cross. This sounds easier than it is, because there are stones blocking the way, slippery ice passages, flipper bumpers and lots more. You can find Mulg at www.ibr.cs.tu-bs.de/~harbaum/pilot/mulg.html.

Ask the Applications Engineer—28

By Eamon Nash


LOGARITHMIC AMPLIFIERS EXPLAINED

Q. I've just been reading data sheets of some recently released Analog Devices log amps and I'm still a little confused about what exactly a log amp does.

A. You're not alone. Over the years, I have had to deal with lots of inquiries about the changing emphasis on functions that log amps perform and radically different design concepts. Let me start by asking you, what do you expect to see at the output of a log amp?

Q. Well, I suppose that I would expect to see an output proportional to the logarithm of the input voltage or current, as you describe in the Nonlinear Circuits Handbook | <http://www.analog.com/publications/magazines/Dialogue/Anniversary/books.html> and the Linear Design Seminar Notes | http://www.analog.com/publications/press/misc/press_123094.html.

A. Well, that's a good start but we need to be more specific. The term log amp, as it is generally understood in communications technology, refers to a device which calculates the log of an input signal's envelope. What does that mean in practice? Take a look at the scope photo below. This shows a 10-MHz sine wave modulated by a 100-kHz triangular wave and the gross logarithmic response of the AD8307, a 500-MHz 90-dB log amp. Note that the input signal on the scope photo consists of many cycles of the 10-MHz signal, compressed together, using the time/div knob of the oscilloscope. We do this to show the envelope of the signal, with its much slower repetition frequency of 100 kHz. As the envelope of the signal increases linearly, we can see the characteristic "log (x)" form in the output response of the log amp. In contrast, if our measurement device were a linear envelope detector (a filtered rectified output), the output would simply be a tri-wave.




Q. So I don't see the log of the instantaneous signal?

A. That's correct, and it's the source of much of the confusion. The log amp gives an indication of the instant-by-instant low-frequency changes in the envelope, or *amplitude*, of the signal in the log domain in the same way that a digital voltmeter, set to "ac volts," gives a steady (linear) reading when the input is connected to a constant amplitude sine wave and follows any

adjustments to the amplitude. A device that calculates the instantaneous log of the input signal is quite different, especially for bipolar signals.

On that point, let's digress for a moment to consider such a device. Think about what would happen when an ac input signal crosses zero and goes negative. Remember, the mathematical function, $\log x$, is undefined for x real and less than or equal to zero, or $-x$ greater than or equal to zero (see figure).



However, as the figure shows, the inverse hyperbolic sine, $\sinh^{-1} x$, which passes symmetrically through zero, is a good approximation to the combination of $\log 2x$ and minus $\log (-2x)$, especially for large values of $|x|$. And yes, it is possible to build such a log amp; in fact, Analog Devices many years ago manufactured and sold Model 752 N & P temperature-compensated log diode modules, which—in complementary feedback pairs—performed that function. Such devices, which calculate the instantaneous log of the input signal are called *baseband log amps* (the term "true log amp" is also used). The focus of this discussion, however, is on envelope-detecting log amps, also referred to as demodulating log amps, which have interesting applications in RF and IF circuitry for communications.


Q. But, from what you have just said, I would imagine that a log amp is generally not used to demodulate signals?

A. Yes, that is correct. The term demodulating came to be applied to this type of device because a log amp recovers the log of the envelope of a signal in a process somewhat like that of demodulating AM signals.

In general, the principal application of log amps is to measure signal *strength*, as opposed to detecting signal *content*. The log amp's output signal, which can represent a many-decade dynamic range of high-frequency input signal amplitudes by a relatively narrow range, is typically used to regulate gain. The classic example of this is using a log amp in an automatic gain control loop, to regulate the gain of a variable-gain amplifier. The receiver of a cellular base station, for example, might use the signal from a log amp to regulate the receiver gain. In transmitters, log amps are also used to measure and regulate transmitted power.

However, there are some applications where a log amp is used to demodulate a signal. The figure shows a received signal that has been modulated using *amplitude shift keying* (ASK). This simple modulation scheme, similar to early transmissions of radar pulses, conveys digital information by transmitting a series of RF bursts (logic 1 = burst, logic 0 = no burst). When this


signal is applied to a log amp, the output is a pulse train which can be applied to a comparator to give a digital output. Notice that the actual amplitude of the burst is of little importance; we only want to detect its presence or absence. Indeed, it is the log amp's ability to convert a signal which varies over a large dynamic range (10 mV to 1 V in this case) into one that varies over a much smaller range (1 V to 3 V) that makes the use of a log amp so appealing in this application.



Q. Can you explain briefly how a log amp works?

A. The figure shows a simplified block diagram of a log amp. The core of the device is a cascaded chain of amplifiers. These amplifiers have linear gain, usually somewhere between 10 dB and 20 dB. For simplicity of explanation, in this example, we have chosen a chain of 5 amplifiers, each with a gain of 20 dB, or 10 \times . Now imagine a small sine wave being fed into the first amplifier in the chain. The first amplifier will amplify the signal by a factor of 10 before it is applied to the second amplifier. So as the signal passes through each subsequent stage, it is amplified by an additional 20 dB.

Now, as the signal progresses down the gain chain, it will at some stage get so big that it will begin to clip (the term *limit* is also used) as shown. In the simplified example, this clipping level (a desired effect) has been set at 1 V peak. The amplifiers in the gain chain would be designed to limit at this same precise level.



After the signal has gone into limiting in one of the stages (this happens at the output of the third stage in the figure), the limited signal continues down the signal chain, clipping at each stage and maintaining its 1 V peak amplitude as it goes.

The signal at the output of each amplifier is also fed into a full wave rectifier. The outputs of these rectifiers are summed together as shown and applied to a low-pass filter, which removes the ripple of the full-wave rectified signal. Note that the contributions of the earliest stages are so small as to be negligible. This yields an output (often referred to as the "video" output), which will be a steady-state quasi-logarithmic dc output for a steady-state ac input signal. The actual devices contain innovations in circuit design that shape the gain and limiting functions to produce smooth and accurate logarithmic behavior between the decade breaks, with the limiter output sum comparable to the *characteristic*, and the contribution of the less-than-limited terms to the *mantissa*.

To understand how this signal transformation yields the log of the input signal's envelope, consider what happens if the input signal is reduced by 20 dB. As it stands in the figure, the unfiltered output of the summer is about 4 V peak (from 3 stages that are limiting and a fourth that is just about to limit). If the input signal is reduced by a factor of 10, the output of one stage at the input end of the chain will become negligible, and there will be one less stage in limiting. Because of the voltage lost from this stage, the summed output will drop to approximately 3 V. If the input signal is reduced by a further 20 dB, the summed output will drop to about 2 V.

So the output is changing by 1 V for each factor-of-10 (20-dB) amplitude change at the input. We can describe the log amp then as having a slope of 50 mV/dB.

Q. O.K. I understand the logarithmic transformation. Now can you explain what the Intercept is?

A. The slope and intercept are the two specifications that define the transfer function of the log amp, that is, the relationship between output voltage and input signal level. The figure shows the transfer function at 900 MHz, and over temperature, of the AD8313, a 100-MHz-to-2.5-GHz 65-dB log amp. You can see that the output voltage changes by about 180 mV for a 10 dB change at the input. From this we can deduce that the slope of the transfer function is 18 mV/dB.


As the input signal drops down below about -65 dBm, the response begins to flatten out at the bottom of the device's range (at around 0.5 V, in this case). However, if the linear part of the transfer function is extrapolated until it crosses the horizontal axis (0 V theoretical output), it passes through a point called the *intercept* (at about -93 dBm in this case). Once the slope and intercept of a particular device are known (these will always be given in the data sheet), we can predict the nominal output voltage of the log amp for any input level within the linear range of the device (about -65 dBm to 0 dBm in this case) using the simple equation:

$$V_{OUT} = \text{Slope} \times (P_{IN} - \text{Intercept})$$

For example, if the input signal is -40 dBm the output voltage will be equal to

$$\begin{aligned} 18 \text{ mV/dB} \times (-40 \text{ dBm} - (-93 \text{ dBm})) \\ = 0.95 \text{ V} \end{aligned}$$

It is worth noting that an increase in the intercept's value *decreases* the output voltage.



The figure also shows plots of deviations from the ideal, i.e., log conformance, at -40°C , $+25^{\circ}\text{C}$, and $+85^{\circ}\text{C}$. For example, at $+25^{\circ}\text{C}$, the log conformance is to within at least ± 1 dB for an input in the range -2 dBm to -67 dBm (over a smaller range, the log conformance is even better). For this reason, we call the AD8313 a 65-dB log amp. We could just as easily say that the AD8313 has a dynamic range of 73 dB for log conformance within 3 dB.

Q. In doing some measurements, I've found that the output level at which the output voltage flattens out is higher than specified in the data sheet. This is costing me dynamic range at the low end. What is causing this?

A. I come across this quite a bit. This is usually caused by the input picking up and measuring an external noise. Remember that our log amps can have an input bandwidth of as much as 2.5 GHz! The log amp does not know the difference between the wanted signal and the noise. This happens quite a lot in laboratory environments, where multiple signal sources may be present. Remember, in the case of a wide-range log amp, a -60 -dBm noise signal, coming from your colleague who is testing his new cellular phone at the next lab bench, can wipe out the bottom 20-dB of your dynamic range.

A good test is to ground both differential inputs of the log amp. Because log amps are generally ac-coupled, you should do this by connecting the inputs to ground through coupling capacitors.

Solving the problem of noise pickup generally requires some kind of filtering. This is also achieved indirectly by using a matching network at the input. A narrow-band matching



network will have a filter characteristic and will also provide some gain for the wanted signal. Matching networks are discussed in more detail in data sheets for the AD8307, AD8309, and AD8313.

Q. What corner frequency is typically chosen for the output stage's low-pass filter?

A. There is a design trade-off here. The corner frequency of the on-chip low-pass filter must be set low enough to adequately remove the ripple of the full-wave rectified signal at the output of the summer. This ripple will be at a frequency 2 times the input signal frequency. However the RC time constant of the low-pass filter determines the maximum rise time of the output. Setting the corner frequency too low will result in the log amp having a sluggish response to a fast-changing input envelope.

The ability of a log amp to respond to fast changing signals is critical in applications where short RF bursts are being detected. In addition to the ASK example discussed earlier, another good example of this is RADAR. The figure on the left shows the response of the AD8313 to a short 100 MHz burst. In general, the log-amp's response time is characterized by the metric 10% to 90% rise time. The table below compares the rise times and other important specifications of different Analog Devices log amps.

Now take a look at the figure on the right. This shows you what will happen if the frequency of the input signal is lower than the corner frequency of the output filter. As might be expected, the full wave rectified signal appears unfiltered at the output. However this situation can easily be improved by adding additional low-pass filtering at the output.



Part Number	Input Bandwidth	10%–90% Rise Time	Dynamic Range	Log Conformance	Limiter Output
AD606	50 MHz	360 ns	80 dB	± 1.5 dB	Yes
AD640	120 MHz	6 ns	50 dB	± 1 dB	Yes
AD641	250 MHz	6 ns	44 dB	± 2 dB	Yes
AD8306	500 MHz	67 ns	95 dB	± 0.4 dB	Yes
AD8307	500 MHz	500 ns	92 dB	± 1 dB	No
AD8309	500 MHz	67 ns	100 dB	± 1 dB	Yes
AD8313	2500 MHz	45 ns	65 dB	± 1 dB	No

Q. I notice that there is an unusual tail on the output signal at the right. What is causing that?


A. That is an interesting effect that results from the nature of the log transformation that is taking place. Looking again at transfer function plot (i.e. voltage out vs. input level), we can see that at low input levels, small changes in the input signal have a significant effect on the output voltage. For example a change in the input level from 7 mV to 700 μ V (or about -30 dBm to -50 dBm) has the same effect as a change in input level from 70 mV to 7 mV. That is what is expected from a logarithmic amplifier. However, looking at the input signal (i.e., the RF burst) with the naked eye, we do not see small changes in the mV range. What's happening in the figure is that the burst does not turn off instantly but drops to some level and then decays exponentially to zero. And the log of a decaying exponential signal is a straight line similar to the tail in the plot.

Q. Is there a way to speed up the rise time of the log amp's output?

A. This is not possible if the internal low-pass filter is buffered, which is the case in most devices. However the figure shows one exception: the un-buffered output stage of the AD8307 is here represented by a current source of 2 μ A/dB, which is looking at an internal load of 12.5 k Ω . The current source and the resistance combine to give a nominal slope of 25 mV/dB. The 5-pF capacitance in parallel with the 12.5-k Ω resistance combines to yield a low-pass corner frequency of 2.5 MHz. The associated 10%-90% rise time is about 500 ns.

In the figure, an external 1.37-k Ω shunt resistor has been added. Now, the overall load resistance is reduced to around 1.25 k Ω . This will decrease the rise time ten-fold. However the overall logarithmic slope has also decreased ten-fold. As a result, external gain is required to get back to a slope of 25 mV/dB.


You may also want to take a look at the Application Note AN-405. This shows how to improve the response time of the AD606.



Q. Returning to the architecture of a typical log amp, is the heavily clipped signal at the end of the gain chain in any way useful?


A. The signal at the end of the linear gain chain has the property that its amplitude is constant for all signal levels within the dynamic range of the log amp. This type of signal is very useful in phase- or frequency demodulation applications. Remember that in a phase-modulation scheme (e.g. QPSK or broadcast FM), there is no useful information contained in the signal's amplitude; all the information is contained in the phase. Indeed, amplitude variations in the signal can make the demodulation process quite a bit more difficult. So the signal at the output of the linear gain chain is often made available to give a limiter output. This signal can then be applied to a phase or frequency demodulator.

The degree to which the phase of the output signal changes as the input level changes is called *phase skew*. Remember, the phase between input and output is generally not important. It is more important to know that the phase from input to output stays constant as the input signal is swept over its dynamic range. The figure shows the phase skew of the AD8309's limiter output, measured at 100 MHz. As you can see, the phase varies by about 6° over the device's dynamic range and over temperature.



Q. I noticed that something strange happens when I drive the log amp with a square wave.

A. Log amps are generally specified for a sine wave input. The effect of differing signal waveforms is to shift the effective value of the log amp's intercept upwards or downwards. Graphically, this looks like a vertical shift in the log amp's transfer function (see figure), without affecting the logarithmic slope. The figure shows the transfer function of the AD8307 when alternately fed by an unmodulated sine wave and by a CDMA channel (9 channels on) of the same rms power. The output voltage will differ by the equivalent of 3.55 dB (88.7 mV) over the complete dynamic range of the device.



The table shows the correction factors that should be applied to measure the rms signal strength of various signal types with a logarithmic amplifier which has been characterized using a sine wave input. So, to measure the rms power of a squarewave, for example, the mV equivalent of the dB value given in the table (−3.01 dB, which corresponds to 75.25 mV in the case of the AD8307) should be subtracted from the output voltage of the log amp.

Signal Type	Correction Factor (Add to Output Reading)
Sine Wave	0 dB
Square Wave or DC	−3.01 dB
Triangular Wave	+0.9 dB
GSM Channel (All Time Slots On)	+0.55 dB
CDMA Forward Link (Nine Channels On)	+3.55 dB
Reverse CDMA Channel	0.5 dB
PDC Channel (All Time Slots On)	+0.58 dB
Gaussian Noise	+2.51 dB

Q. In your data sheets you sometimes give input levels in dBm and sometimes in dBV. Can you explain why?


A. Signal levels in communications applications are usually specified in dBm. The dBm unit is defined as the power in dB relative to 1 mW i.e.,

$$\text{Power (dBm)} = 10 \log_{10} (\text{Power}/1 \text{ mW})$$

Since power in watts is equal to the rms voltage squared, divided by the load impedance, we can also write this as

$$\text{Power (dBm)} = 10 \log_{10} ((V_{\text{rms}}^2/R)/1 \text{ mW})$$

It follows that 0 dBm occurs at 1 mW, 10 dBm corresponds to 10 mW, +30 dBm corresponds to 1 W, etc. Because impedance is a component of this equation, it is always necessary to specify load impedance when talking about dBm levels.



Log amps, however fundamentally respond to voltage, not to power. The input to a log amp is usually terminated with an


external 50-Ω resistor to give an overall input impedance of approximately 50 Ω, as shown in the figure (the log amp has a relatively high input impedance, typically in the 300 Ω to 1000 Ω range). If the log amp is driven with a 200-Ω signal and the input is terminated in 200 Ω, the output voltage of the log amp will be higher compared to the same amount of power from a 50-Ω input signal. As a result, it is more useful to work with the *voltage* at the log amp's input. An appropriate unit, therefore, would be dBV, defined as the voltage level in dB relative to 1 V, i.e.,

$$\text{Voltage (dBV)} = 20 \log_{10} (V_{\text{rms}}/1 \text{ V})$$

However, there is disagreement in the industry as to whether the 1-V reference is 1 V peak (i.e., amplitude) or 1 V rms. Most lab instruments (e.g., signal generators, spectrum analyzers) use 1 V rms as their reference. Based upon this, dBV readings are converted to dBm by adding 13 dB. So −13 dBV is equal to 0 dBm.

As a practical matter, the industry will continue to talk about input levels to log amps in terms of dBm power levels, with the implicit assumption that it is based on a 50 Ω impedance, even if it is not completely correct to do so. As a result it is prudent to provide specifications in *both* dBm and dBV in data sheets.

The figure shows how mV, dBV, dBm and mW relate to each other for a load impedance of 50 Ω. If the load impedance were 20 Ω, for example, the V (rms), V (p-p) and dBV scales will be shifted downward relative to the dBm and mW scales. Also, the V (p-p) scale will shift relative to the V (rms) scale if the peak to rms ratio (also called crest factor) is something other than $\sqrt{2}$ (the peak to rms ratio of a sine wave). ■



Analog-to-Digital Converter Architectures and Choices for System Design

By Brian Black

How important are the differences between sigma-delta and successive-approximation architectures in choosing an analog-to-digital (A/D) converter? They can often be an important factor in initiating the selection of a converter for a specific application. We describe here four major circuit architectures used in A/D converter (ADC) design and outline the role they play in converter choice for various kinds of applications. The descriptions are augmented by three examples that illustrate tradeoffs and issues associated with architectural considerations.

Though not detailed or exhaustive, this overview is intended to raise issues that should be understood when considering converters of different architectures. Sources of more-detailed information on converter architectures can be found in the *References* and at Internet sites indicated at appropriate points. As one might expect in a survey of this kind, these descriptions are not comprehensive; and variations within each of the architecture families make generalizations less than fully accurate. Nevertheless, such generalizations are useful for the system designer to keep in mind when conducting a high level overview of a proposed system's requirements.

CONVERTER ARCHITECTURES

An overwhelming variety of ADCs exist on the market today, with differing resolutions, bandwidths, accuracies, architectures, packaging, power requirements, and temperature ranges, as well as hosts of specifications, covering a broad range of performance needs. And indeed, there exists a variety of applications in data-acquisition, communications, instrumentation, and interfacing for signal processing, all having a host of differing requirements.

Considering architectures, for some applications just about any architecture could work well; for others, there is a "best choice." In some cases the choice is simple because there is a clear-cut advantage to using one architecture over another. For example, pipelined converters are most popular for applications requiring a throughput rate of more than 5 MSPS with good resolution. Sigma-delta converters are usually the best choice when very high resolution (20 bits or more) is needed. But in some cases the choice is more subtle. For example, the sigma-delta AD7722 and the successive-approximations AD974 have similar resolution (16 bits) and throughput performance (200 kSPS). Yet the differences in their underlying architectures make one or the other a better choice, depending on the application.

The most popular ADC architectures available today are *successive approximations* (sometimes called SAR because a *successive-approximations (shift) register* is the key defining element), *flash* (all decisions made simultaneously), *pipelined* (with multiple flash stages), and *sigma-delta* ($\Sigma\Delta$), a charge-balancing type. All A/D converters require one or more steps involving comparison of an input signal with a reference. Figure 1 shows qualitatively how

flash, pipelined, and SAR architectures differ with respect to the number of comparators used vs. the number of comparison cycles needed to perform a conversion.

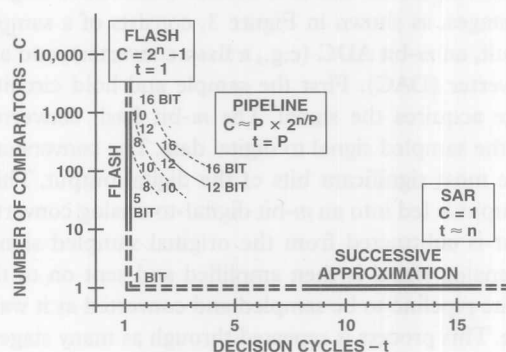


Figure 1. Tradeoff between decision cycles and comparators.

FLASH CONVERTERS

Conceptually, the *flash* architecture (illustrated in Figure 2) is quite straightforward: a set of $2^n - 1$ comparators is used to directly measure an analog signal to a resolution of n bits. For a 4-bit flash ADC, the analog input is fed into 15 comparators, each of which is biased to compare the input to a discrete transition value. These values are spaced one least-significant bit (LSB = $FS/2^n$) apart. The comparator outputs simultaneously present $2^n - 1$ discrete digital output states. If for example the input is just above 1/4 of full scale, all comparators biased to less than 1/4 full scale will output a digital "1," and the others will output a digital "0." Together, these outputs can be read much like a liquid thermometer. The final step is to level-decode the result into binary form.

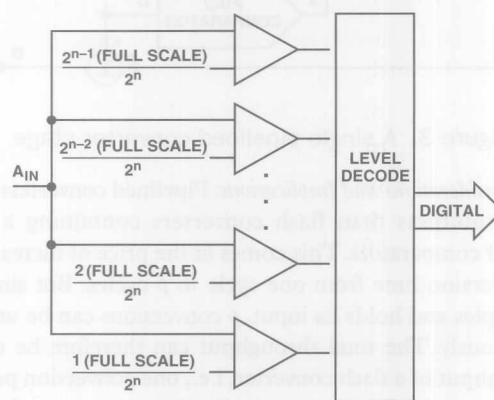


Figure 2. Basic flash architecture.

Design Considerations and Implications: The flash architecture has the advantage of being very fast, because the conversion occurs in a single ADC cycle. The disadvantage of this approach is that it requires a large number of comparators that are carefully matched and properly biased to ensure that the results are linear. Since the number of comparators needed for an n -bit resolution ADC is equal to $2^n - 1$, limits of physical integration and input loading keep the maximum resolution fairly low. For example, a 4-bit ADC requires 15 comparators, an 8-bit ADC requires 255 comparators, and a 16-bit ADC would require 65,535 comparators! For more about flash ADCs, see http://www.analog.com/support_standard_linear/seminar_material/practical_design_techniques/Section4.pdf.

PIPELINED ARCHITECTURE

The *pipelined* (or pipelined-flash) architecture effectively overcomes the limitations of the flash architecture. A pipelined converter divides the conversion task into several consecutive stages. Each of these stages, as shown in Figure 3, consists of a sample-and-hold circuit, an m -bit ADC (e.g., a flash converter), and an m -bit D/A converter (DAC). First the sample and hold circuit of the first stage acquires the signal. The m -bit flash converter then converts the sampled signal to digital data. The conversion result forms the most significant bits of the digital output. This same digital output is fed into an m -bit digital-to-analog converter, and its output is subtracted from the original sampled signal. The residual analog signal is then amplified and sent on to the next stage in the pipeline to be sampled and converted as it was in the first stage. This process is repeated through as many stages as are necessary to achieve the desired resolution. In principle, a pipelined converter with p pipeline stages, each with an m -bit flash converter, can produce a high-speed ADC with a resolution of $n = p \times m$ bits using $p \times (2^m - 1)$ comparators. For example, a 2-stage pipelined converter with 8-bit resolution requires 30 comparators, and a 4-stage 16-bit ADC requires only 60 comparators. In practice, however, a few additional bits are generated to provide for error correction. For more about pipelined ADCs, See http://www.analog.com/support/standard_linear/seminar_material/practical_design_techniques/Section4.pdf.

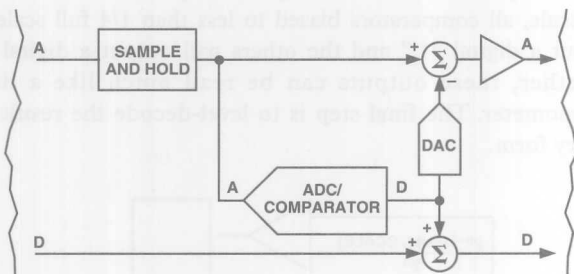


Figure 3. A single pipelined converter stage.

Design Considerations and Implications: Pipelined converters achieve higher resolutions than flash converters containing a similar number of comparators. This comes at the price of increasing the total conversion time from one cycle to p cycles. But since each stage samples and holds its input, p conversions can be underway simultaneously. The total throughput can therefore be equal to the throughput of a flash converter, i.e., one conversion per cycle. The difference is that for the pipelined converter, we have now introduced *latency* equal to p cycles. Another limitation of the pipelined architecture is that the conversion process generally requires a clock with a fixed period. Converting rapidly varying non-periodic signals on a traditional pipelined converter can be difficult because the pipeline typically runs at a periodic rate.

SUCCESSIVE APPROXIMATIONS

The successive-approximations architecture can be thought of as being orthogonal to the flash architecture. While a flash converter uses many comparators to convert in a single cycle; a SAR converter, shown in Figure 4, conceptually uses a single comparator over many cycles to make its conversion. The SAR converter works like an old-fashioned balance scale. On one side of the scale, we place the sampled unknown quantity. On the other side, we place a weight (generated by the SAR and DAC) that has the value of $1/$

2 of full-scale and compare the two values. This first weight represents the most significant bit (MSB). If the unknown quantity is larger, the $1/2$ -scale weight is retained; if the unknown quantity is smaller, it is removed. This series of steps is repeated n times, using successively smaller weights in binary progression (e.g., $1/4$, $1/8$, $1/16$, $1/32$, $\dots 1/2^n$ of full scale) until the desired resolution, n , is attained. Each weight represents a binary bit, with the largest representing the most significant bit, and the smallest representing the least significant bit.




Figure 4. Successive-approximations architecture.

Design Considerations and Implications: A SAR converter can use a single comparator to realize a high resolution ADC. But it requires n comparison cycles to achieve n -bit resolution, compared to p cycles for a pipelined converter and 1 cycle for a flash converter. Since a successive-approximations converter uses a fairly simple architecture employing a single SAR, comparator, and DAC, and the conversion is not complete until all weights have been tested, only one conversion is processed during n comparison cycles. For this reason, SAR converters are more often used at lower speeds in higher-resolution applications. SAR converters are also well suited for applications that have non-periodic inputs, since conversions can be started at will. This feature makes the SAR architecture ideal for converting a series of time-independent signals. A single SAR converter and an input multiplexer are typically less expensive to implement than several sigma-delta converters. With dither noise present, SAR and pipelined converters can use averaging to increase the effective resolution of the converter: for every doubling of sample rate, the effective resolution improves by 3 dB or $1/2$ bit.

One consideration when using a SAR or pipelined converter is *aliasing*. The process of sampling a signal leads to aliasing—the frequency-domain reflection of signals about the sampling frequency. In most applications, aliasing is an unwanted effect that requires a low-pass anti-alias filter ahead of the ADC to remove high-frequency noise components, which would be aliased into the passband. However, *undersampling* can put aliasing to good use, most often in communications applications, to convert a high-frequency signal to a lower frequency. Undersampling is effective as long as the total bandwidth of a signal meets the Nyquist criterion (less than one-half the sampling rate), and the converter has sufficient acquisition and signal sampling performance at the higher frequencies where the signal resides. While fast SAR converters are capable of undersampling, the faster pipelined converters tend to be more effective at it. For more about undersampling and dither, see http://www.analog.com/support/standard_linear/seminar_material/practical_design_techniques/Section5.pdf.

SIGMA-DELTA

The sigma-delta architecture takes a fundamentally different approach than those outlined above. In its most basic form, a sigma-delta converter consists of an integrator, a comparator, and a single-bit DAC, as shown in Figure 5. The output of the DAC is subtracted from the input signal. The resulting signal is then integrated, and

the integrator output voltage is converted to a single-bit digital output (1 or 0) by the comparator. The resulting bit becomes the input to the DAC, and the DAC's output is subtracted from the ADC input signal, etc. This closed-loop process is carried out at a very high "oversampled" rate. The digital data coming from the ADC is a stream of "ones" and "zeros," and the value of the signal is proportional to the density of digital "ones" coming from the comparator. This bit stream data is then digitally filtered and decimated to result in a binary-format output. For more about sigma-delta conversion, see http://www.analog.com/support/standard_linear/seminar_material/practical_design_techniques/Section3.pdf.




Figure 5. Sigma-delta ADC architecture.

Design Considerations and Implications: One of the most advantageous features of the sigma-delta architecture is the capability of *noise shaping*, a phenomenon by which much of the low-frequency noise is effectively pushed up to higher frequencies and out of the band of interest. As a result, the sigma-delta architecture has been very popular for designing low-bandwidth high-resolution ADCs for precision measurement. Also, since the input is sampled at a high "oversampled" rate, unlike the other architectures described in this paper, the requirement for external anti-alias filtering is greatly relaxed. A limitation of this architecture is its *latency*, which is substantially greater than that of the other types. Because of oversampling and latency, sigma-delta converters are not often used in multiplexed signal applications. To avoid interference between multiplexed signals, a delay at least equal to the decimator's total delay must occur between conversions. These characteristics can be improved in sophisticated sigma-delta ADC designs by using multiple integrator stages and/or multi-bit DACs.

APPLICATION EXAMPLES

The following three examples illustrate some of the issues described above.

Example 1: Multiple Inputs, 16-Bit Resolution

Consider an application that requires 16-bit resolution for 4 independent signals with bandwidths of dc to 15 kHz, 15 kHz, 15 kHz, and 45 kHz. The total throughput required to sample these signals under the Nyquist criterion is $(2 \times 15 + 2 \times 15 + 2 \times 15 + 2 \times 45)$ kSPS = 180 kSPS. At first glance, the SAR-type AD974, the sigma-delta AD7722, and the pipelined/sigma-delta AD9260 all have the required total throughput capability. But, as has been discussed above, the inherent latency of sigma-delta converters limits their effective throughput when they must continually acquire new signals by multiplexing. *Effective multiplexed throughput* can be defined as the total throughput of a converter when two or more independent signals are multiplexed. The following table compares the total throughput and effective throughput for each converter and indicates the number of converters of its type that would be needed to serve in this application.

Converter	Architecture	Total Throughput (16-Bit Resolution)	Effective Multiplexed Throughput	Converters Needed for Application
AD974	SAR	200 kSPS	200 kSPS	1
AD7722	Sigma-Delta	195 kSPS	2.3 kSPS	4
AD9260	Pipelined/Sigma-Delta	2500 kSPS	75 kSPS	3

Whether converting a single input or several multiplexed inputs, the AD974 achieves a throughput rate of up to 200 kSPS. Since the application requires a total throughput of 180 kSPS, the AD974's performance is sufficient. In fact, this is exactly the type of application that the AD974 was designed for: in addition to the SAR converter and reference, it also contains an integrated 4-channel multiplexer.

The AD7722 and AD9260 both face the challenges confronted by sigma-delta converters in multiplexing several inputs. The AD7722's throughput is 195 kSPS when sampling a single signal, but it drops to just 2.3 kSPS when converting multiple signals, due to the settling time which results from oversampling and filtering. To use the AD7723 in this application, four converters (one per channel) would be needed.

The AD9260 combines pipeline and sigma-delta techniques. Its throughput rate of 2.5 MSPS makes it ideal for higher throughput single channel systems. But in this application, its settling time of 13.35 ms limits its effective throughput to 75 kSPS. To use the AD9260 in this application would require at least 3 converters. Note that if the AD9260 were purely a pipelined flash converter, a single converter would have had the required throughput, assuming that the inputs are periodic.

Example 2: Single Input, 16-Bit Resolution

Consider now an application that converts a single 90-kHz bandwidth input at 16-bit resolution. In this case, all three converters from the first example would work well. Here the choice among converters would be made on other considerations, including ac and dc performance, system-level considerations (e.g., Is there a great benefit to the anti-alias performance of sigma-delta converters in this application?), latency, and cost.

Example 3: Multiple Inputs, 14-Bit Resolution

Consider an application in which 16 inputs, each with a dc to 100-kHz bandwidth, are converted with a resolution of at least 14 bits. Three converters suitable for this application include the SAR-type AD7865, the sigma-delta AD7722, and the pipelined AD9240. The total throughput required under the Nyquist criterion is $2 \times 100 \text{ kHz} \times 16 = 3.2 \text{ MSPS}$. The following table shows the throughput for each converter.

Converter	Architecture	Total Throughput (14-Bit Resolution)	Effective Multiplexed Throughput	Converters Needed for Application
AD7865	SAR	416 kSPS	416 kSPS	8
AD7722	Sigma-Delta	220 kSPS	2.3 kSPS	16
AD9240	Pipelined	10 MSPS	10 MSPS	1

Of the three converters, only the AD9240 has the throughput needed to convert all 16 channels. The AD7865 has sufficient throughput for 2 inputs per converter. To use the AD7865 in this application, 8 converters would be needed. The AD7722 would

need to be used in a converter-per-channel implementation; thus 16 converters would be required.

Summary

While not exhaustive, the following table summarizes and ranks (in a generalized sense) the relative advantages of flash, pipelined, SAR, and sigma-delta architectures. A rank of 1 in a performance category indicates that the architecture is inherently better than the others for that category. An * indicates that the architecture has the capability or characteristic listed.

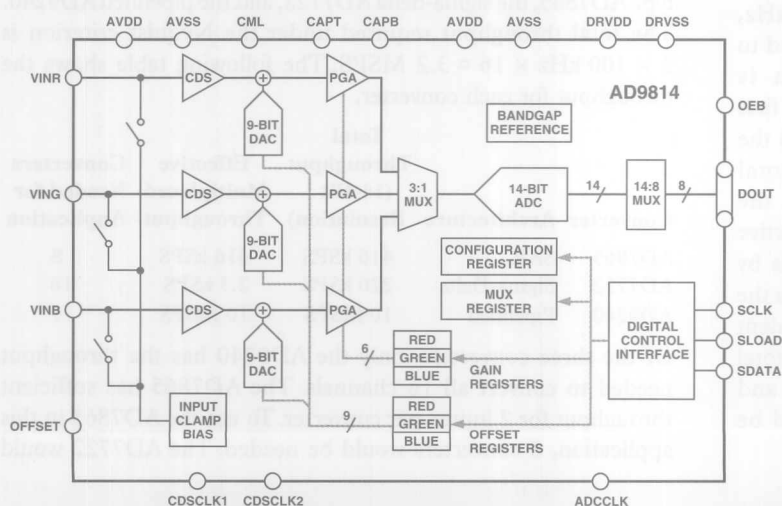
Characteristic	Flash	Pipelined	SAR	Sigma-Delta
Throughput	1	2	3	4
Resolution (ENOB)	4	3	2	1
Latency	1	3	2	4
Suitability for converting multiple signals per ADC	1	2	1	3
Capability to convert non-periodic multiplexed signals	1	2	1	3
Simplified anti-aliasing				*
Can undersample	*	*	*	
Can increase resolution through averaging (with dither noise)	*	*	*	

New-Product Briefs

CCD SIGNAL PROCESSOR

The first true 14-bit CCD processor with 3-channel input and digitizer

The AD9814 is a 14-bit, 3-channel analog signal processor and digitizer designed for high-performance 3-channel 10 MSPS tri-linear CCD imaging applications, such as color film- and document scanners. Each channel includes an input clamp, correlated double sampler (CDS), 9-bit offset DAC, and 6-bit programmable-gain amplifier (PGA), multiplexer, and 14-bit ADC. The output is a pair of 8-bit multiplexed words, accessed by two Read cycles. The device is programmed via a serial interface. The true-14-bit K grade guarantees no missing codes, maximum differential nonlinearity of ± 1 LSB, and the lowest noise in the industry. The device is housed in a 28-lead 0.3" SOIC, operates from 0°C to +70°C, and requires a 3.0 to 5.25 V supply. Prices in 100s (K/J versions) are \$29.50 and \$11.15.



ACKNOWLEDGMENT


The author is indebted to Alain Guery and Mike Coln for providing valuable direction, and for providing the figures shown.

REFERENCES

Analog-Digital Conversion Handbook, The Engineering Staff of Analog Devices, Inc. Englewood Cliffs, NJ: Prentice Hall, 1986. ISBN 0-13-032848-0. Available from Analog Devices: Phone (781) 461-3392.

Practical Analog Design Techniques, edited by Walt Kester. Norwood, MA: Analog Devices, Inc., 1995. ISBN 0-916550-16-8. Available from Analog Devices: Phone (781) 461-3392. Also available free on the Internet as *PDF chapters*. http://www.analog.com/support/standard_linear/seminar_material/index.html

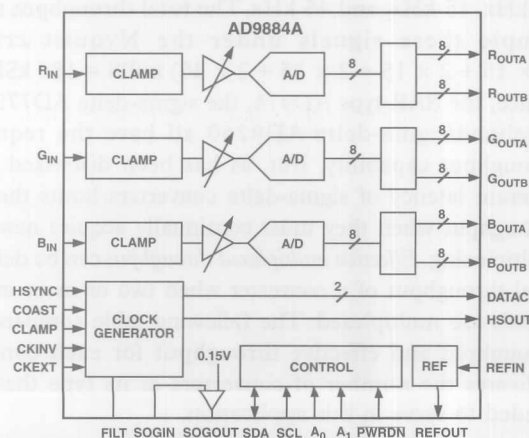
High Speed Design Techniques, edited by Walt Kester. Norwood, MA: Analog Devices, Inc., 1996. ISBN 0-916550-17-6. Available from Analog Devices: Phone (781) 461-3392. Also available free on the Internet as *PDF chapters*. http://www.analog.com/support/standard_linear/seminar_material/index.html

Linear Design Seminar, edited by Walt Kester. Norwood, MA: Analog Devices, Inc., 1995. ISBN 0-916550-15-X. Available from Analog Devices: Phone (781) 461-3392. 

AFE FOR FLAT-PANEL DISPLAYS

First Integrated Analog Interface for SXGA Flat Panel Displays (100 MSPS/140 MSPS)

The AD9884A is the World's first integrated analog interface for SXGA flat-panel displays. Integrating a triple 140 MSPS A/D converter, PLL, and preamplifiers, it is a monolithic analog interface optimized for digitizing RGB graphics signals from personal computers and workstations. Its 140 MSPS encode-rate capability, with 500 MHz full-power analog bandwidth, supports display resolutions of up to 1280×1024 at 75 Hz with sufficient bandwidth to accurately acquire and digitize each pixel. Operating on a 3.3 V supply, it dissipates less than 650 mW. Its signal-to-noise ratio is better than 46 dB at 20 MHz inputs. Operating from 0 to +70°C, it is available in 128-pin MQFP packages in two sampling-rate grades, 100 MSPS (XGA) and 140 MSPS (SXGA), priced at \$21 and \$25.50 in 1000s. An evaluation board is available.



New TxDAC[®] Generation

125-MSPS 10-, 12-, and 14-bit high-performance DACs for wideband multitone communication Transmit channels

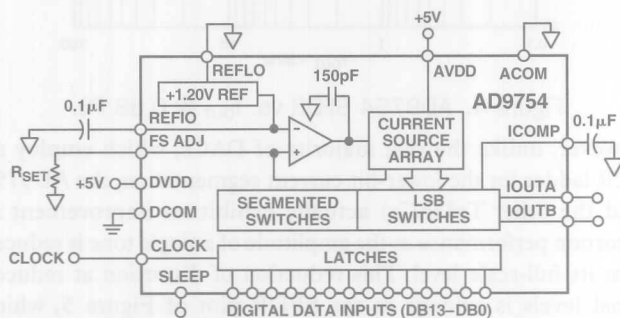


Figure 1. AD9754 functional block diagram.

The AD9754, AD9752, and AD9750 are 14-, 12-, and 10-bit wideband, high-performance, low-power CMOS digital-to-analog converters that operate from +5V supplies. They offer exceptional ac and dc performance, while supporting update rates up to 125 MSPS. All three share the same interface options, 28-lead TSSOP or 0.3" small-outline package, and pinout as the first generation, the AD976x series [Analog Dialogue 30-3, 1996, pp. 3-5]. This allows a highly flexible upward or downward selection path, based on performance, resolution, and cost.

These DACs are current-output; they have a nominal 20-mA full-scale output with a 100-k Ω source impedance and 1.25-V compliance range. The outputs are differential to support either differential or single-ended applications. The devices include an internal 1.20-V reference and a control amplifier which can set output full-scale values from 2 mA for power conservation to 20 mA, providing 20 dB of output flexibility. They have low dissipation of 190 mW, which can be reduced to 65 mW at the 2-mA end of the output scale, and just 20 mW in power-down mode. They are equipped with edge-triggered latches, and the digital interface is CMOS compatible (+2.7 to 5.5 V).

Emerging wireless/wireline communications standards (Wideband, Multicarrier) are demanding lower distortion and noise to improve system capacity and signal quality. The D/A converter in a modem/transceiver transmitting chain (TxDAC) is the basic analog signal generator; it defines the ultimate performance available in a communications system. Everything after it—power amplifier, antenna, transmission medium, and receiver front end—can only contribute to the degradation of a digitally defined signal, and the

end result can never be better than what is at first available from the TxDAC.

The AD975x family is specifically designed for applications such as wideband CDMA, software/multi-carrier base stations, wireless local-loop RIUs, wireless LAN, broadband set-top boxes, and xDSL modems that utilize wide signal bandwidths to transmit high volumes of data over a desired medium. Key specifications over a 25-MHz band, such as 68-dBc SFDR (spurious-free dynamic range) and -65-dBc THD (total harmonic distortion), are improved over those of the AD976x family by 8 to 10 dB, accompanied by a 9-dB lower integrated noise floor at -109 dB.

The AD975x DACs operate over the industrial temperature range from -40 to +85°C. Evaluation boards are available. Prices* (1000s) for 14/12/10 bits are \$21.09/\$16.70/\$9.45.

What's special about DACs for multitone?

The frequency-domain performance of high-speed DACs has traditionally been characterized by analyzing the spectral output of a reconstructed full-scale (i.e., 0 dB FS) single-tone sine wave at a particular output frequency and update rate. Although this characterization data is useful, it is often insufficient to reflect a DAC's performance for a reconstructed multitone or spread-spectrum waveform. In fact, evaluating a DAC's spectral performance using a full-scale single-tone at the highest specified frequency (i.e., f_H) of a bandlimited waveform is typically indicative of a DAC's "worst-case" performance for that given waveform. In the time domain, the full-scale sine wave represents the lowest peak-to-rms crest factor (i.e., V_{PEAK}/V_{RMS}) that the bandlimited signal will encounter.

However, the inherent nature of a multitone, spread-spectrum, or QAM waveform, in which the spectral energy of the waveform is spread over a designated bandwidth, will result in a higher peak-to-rms ratio when compared to the case of a simple sine wave. As the reconstructed waveform's peak-to-average ratio increases, an increasing amount of the signal energy is concentrated around the DAC's midscale value. Figure 2 is just one example of a bandlimited multitone vector (i.e., eight tones) centered on one-half the Nyquist bandwidth (i.e., $f_{CLOCK}/4$).




Figure 2. Multitone spectral plot.

This particular multitone vector has a peak-to-rms ratio of 13.5 dB, as compared to a sine wave's peak-to-rms ratio of 3 dB. This means that the DAC must be able to handle a waveform where most of

*Prices indicated here are recommended resale prices (U.S. Dollars) FOB U.S.A. Prices are subject to change without notice. For specific price quotations, get in touch with our sales offices or distributors.

the information is carried at low levels in order to be able to handle large instantaneous peaks without clipping. A "snapshot" of this reconstructed multitone vector in the time domain as shown in Figure 3 reveals the dense signal content around the midscale value. As a result, a DAC's "small-scale" dynamic and static linearity become increasingly critical to obtaining low intermodulation distortion and maintaining sufficient carrier-to-noise ratios for a given modulation scheme.




Figure 3. Time-domain "snapshot" of the multitone waveform.

A DAC's small-scale linearity performance is also an important consideration in applications where additive dynamic range is required for gain-control purposes or "predistortion" signal conditioning. For instance, a DAC with sufficient dynamic range can be used to provide additional digitally controlled variation of its reconstructed signal. In fact, the gain can be controlled in 6-dB increments by simply performing a shift-left or -right on the DAC's digital input word. An application might be to predistort the DAC's input signal intentionally to compensate for nonlinearities associated with subsequent analog components in the signal chain. For example, the signal compression associated with a power amplifier can be compensated for by predistorting the DAC's digital input with the inverse nonlinear transfer function of the power amplifier. Since the DAC must accommodate increased gains at higher output levels, the DAC's performance at reduced signal levels should be carefully evaluated.

A full-scale single tone will induce all of the dynamic and static nonlinearities present in a DAC that contribute to its distortion; the distortion components will reduce its SFDR (Spurious-free dynamic range). Thus, in Figure 4 as the frequency of the digitally generated single-tone waveform increases, the dynamic linearities

of a DAC (here the AD9754) tend to dominate, thus contributing to the substantial reduction in SFDR with increased frequency.




Figure 4. AD9754 SFDR vs. f_{OUT} @ 0 dB FS.

However, unlike the vast majority of DACs, which employ an R-2R ladder for the lower-bit current segmentation, the AD9754 (and the other TxDACs) actually exhibit an improvement in distortion performance as the amplitude of a single tone is reduced from its full-scale level. This reduction of distortion at reduced signal levels is evident in the SFDR plot of Figure 5, which shows typical SFDR vs. frequency at 0, -6, and -12 dB levels for a 65-MSPS sampling rate. Maintaining such decent "small-scale" linearity across the full span of a DAC transfer function is critical in maintaining excellent multitone performance. ▶




Figure 5. AD9754 SFDR vs. f_{OUT} at 65 MSPS.

New-Product Brief

D/A CONVERTER

14-Bit, 150 MSPS TxDAC+® with 2× Interpolation Filter

The AD9772 is the first DAC to be able to do direct IF synthesis in the 75–140 MHz frequency range. It is a single-supply, 14-bit D/A converter with differential current outputs, optimized for baseband or IF waveform reconstruction applications requiring exceptional dynamic range. It is designed for communication Transmit channels, such as WCDMA base stations, multicarrier base stations, and direct IF synthesis. It integrates a complete,

low-distortion 14-bit DAC with a 2× interpolation filter and a PLL clock multiplier with a flexible differential clock-driver input for single-ended or differential clock drive and optimum jitter performance. The AD9772 can reconstruct full-scale waveforms with bandwidths as high as 63.3 MHz while operating at an input data rate of 150 MSPS. It operates from a single 2.7 V to 3.6 V supply, dissipates only 205 mW (13 mW in Power-down), is available in a 48-lead LQFP package, and operates from -40 to +85°C. Price is \$32.18 in 1000s.

Ultrasound Analog Electronics Primer

by Bill Odom

INTRODUCTION

Modern technology is driving medical ultrasound machines to new heights of performance, resulting in images with new levels of clarity and resolution. Ultrasound is used for imaging in cardiac, obstetric, and many other diagnostic areas. Even as the abilities of the machines increase, the cost of machines is coming down. Although ultrasound relies substantially on digital processing, the key to its performance lies in a heavily analog technology. We will examine here the contributions of the analog and mixed-signal components used in ultrasound imaging. Naturally, since each element would by itself need a chapter or a book of its own to provide fully useful information for the system designer, this article is intended to provide an overview and furnish a basic understanding of medical ultrasound architecture. Though this article will barely scratch the surface of design objectives and rules; some of the issues to be raised will have relevance in a broader context of applications and are likely to be examined in greater detail in future articles.

Obtaining an Image

Images are obtained by sweeping a narrow beam of acoustic energy into a living body and analyzing the pattern of the energy reflected back by structures within the body, much like a search radar. Since the receiving transducers are dealing with analog signals, but the analysis is performed digitally, the signals must be digitized. Electrical pulses are applied to piezoelectric ceramic elements to generate the energy at transmitted frequencies from 2 to 20 MHz. The frequency used depends on the application. Higher frequencies provide the best resolution but have less penetration, since they attenuate faster as the signals move through the body. There are limits on how strong the high-frequency pulse can be, because excessive amounts of power are unhealthy for the patient. The most commonly used frequencies range from 2 to 7 MHz.

Return levels range from 1-V echoes near the surface of the body to less than 10 μ V for images deep within the body. The signals are conducted to and from the ceramic elements in the hand piece to the front-end electronics via a cable, which will be subject to noise and attenuation. The wide range of signals must be amplified to 2 V for driving an analog-to-digital converter. To accomplish this, a *time-gain compensation (TGC) amplifier* is used. It will compensate for exponential signal decay by amplifying the signal by an exponential factor that depends on how long the machine has been waiting for the return pulse.

Power levels, frequencies used, amplification, and beam focus determine the clarity of the image. These things are controlled by the sonographer (technician), interacting with the system's inherent properties.

Imaging Modes Used

1. *Gray scale*—produces a basic black-and-white image. It will resolve artifacts as small as 1 mm. The display is made by transmitting bursts of energy and analyzing the return energy (as mentioned above).
2. *Doppler*—The best analogy to medical Doppler ultrasound is color Doppler weather radar. As the name suggests, Doppler

modes detect the velocity of an object in motion by tracking the frequency shift of the return signal. These principles are applied in examining blood or other fluids flowing within the body. It is accomplished by transmitting a continuous wave into the body and producing a fast Fourier transform (FFT) of the return. The computational process will determine the frequency components of the signals from the body and their relationship as a function of fluid velocity. One bin will contain the fundamental transmitted frequency while other bins contain the Doppler shifted frequencies. 4 \times oversampling is often employed.

3. *Venous and arterial modes*—They employ Doppler in conjunction with the gray-scale mode. First the image of a vein or artery will be found. The operator will dial in a small cursor window around it. The Doppler is then engaged within the cursor area. The transmitted signal's Doppler frequency shift will be measured as discussed above. Audio will also be used with the cursor image. Venous flow produces a rushing sound (like a waterfall), while the thump of a pulse will indicate arterial flow. At the same time, blood velocity will be displayed on a digital readout. The sinus rhythm will be displayed as an X-Y plot on the screen. The velocity and rhythm displays are obtained by processing the audio signal from the Doppler shifts.

The Overall System

The block diagram (Figure 1) shows the elements of a system: transducer, multiplexer, transmitter and its beam-forming apparatus, transmit/receive (T/R) switches, low-noise amplifier, signal- and image-processing display, audio, A/D converter and its driver, the TGC amplifier. At the current state of the art, machines can employ as many as 256 channels (comprising 256 ceramic elements, amplifiers, ADCs, etc.).

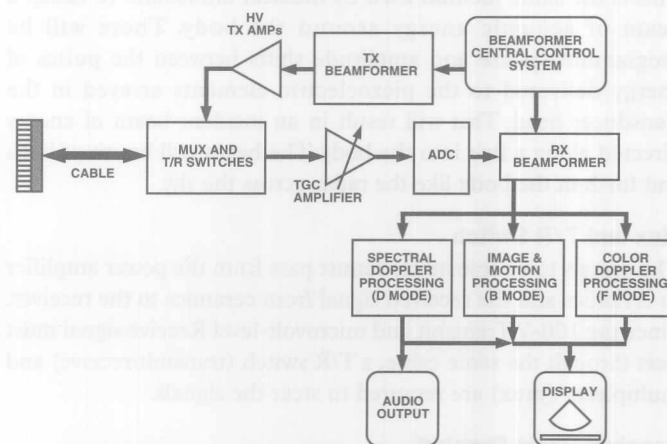


Figure 1. System block diagram.

Probes and Their Transmitted Signals

The probes will have a ceramic element for each channel (up to 256). The elements are made of a piezoelectric ceramic material such as lead zirconium titanate.

In some designs the pulses ring in bursts of a few cycles each time they get a short transmit pulse of about 100 ns ("ping and ring"). The excitation pulse amplitudes will be of the order of 100 V. The magnitude of the pulse will determine the amount of energy beamed into the patient.

In order to minimize distortion, some systems transmit a Gaussian pulse. Figure 2 contrasts the distorted spectrum of a broadband pulse after it is bounced around in the body. Its spectrum bears little resemblance to that of the transmitted pulse. The misshapen pulse will show harmonic distortion and unwanted spurious artifacts. On the other hand, the response to a transmitted Gaussian pulse spectrum looks much the same as when it went out, free of side lobes.




Figure 2. Waveform types (simplified): Gaussian burst, reflected flat-top pulse, and reflected Gaussian pulse.

The excitation pulse might be the output of a DAC, with the signal given the desired shape by a digital synthesizer. The low amplitude pulse will then be amplified to the required amplitude (approximately 100 V).

The receiver must have a wide bandwidth to accommodate the wide range of complex frequencies that must be handled in the DSP's FFT computation. Fast edge rates increase the demand for bandwidth.

Beam Steering and Focus

In the old days of radar, a dish or a banana-shaped antenna would rotate, looking for targets in all directions. As it slowly swept around, a magnetron would fire pulses of energy into the sky. Traveling at the speed of light, reflected energy would come back to the receiver before the antenna moved out of sync. Nowadays, the rotation is produced by phased arrays. The beam is manipulated by varying the phase and power of the signal between antenna radiators, and the beam is swept around the sky without any moving parts.

This is the same method used by medical ultrasound to sweep a beam of acoustic energy around the body. There will be programmed phase and amplitude shifts between the pulses of energy delivered to the piezoelectric elements arrayed in the transducer head. This will result in an incident beam of energy directed along a line into the body. The beam will be swept back and forth in the body like the radar across the sky.

Mux and T/R Switch

The signals to be transmitted must pass from the power amplifier to ceramics and the received signal from ceramics to the receiver. Since the 100-V Transmit and microvolt-level Receive signal must pass through the same cable, a T/R switch (transmit/receive) and multiplexer (mux) are required to steer the signals.

Receiver Beam Forming

The beam is focused by delaying each of the channels so that the return pulses from the *focal point (or area)* arrive at the processor at the same time (see Figure 3). The machine will establish the focal area as set by the operator. Beam forming is currently done with both analog and digital techniques. The machine will adjust the delay required for focus in calculating the position of the sweep line. It will compute the corresponding pixels of the display by using the delay required by each channel to focus the image. Newer machines have multiple focus zones.

Time Gain Control (TGC)

The TGC (time gain compensation) amplifier is a crucial link in the ultrasound signal path. It must have the ability to amplify signals




Figure 3. Beam focus using variable delay.

ranging from a few microvolts to 1 volt up to one or two volts for the ADC. This gain will be exponentially increased along each transmit/receive sweep line. At the near end of the wedge, the gain will be very low. It will have to process the 1-V return signal right after the 100-V ceramic excitation pulse. As time after the excitation pulse passes, the gain will be swept into very high levels. This must be done while maintaining very low noise to avoid masking low-level signal coming from deep within the body. The operator will adjust the TGC amplifier control to improve the quality of the image. The AD604 <http://products.analog.com/products/info.asp?product=AD604> variable-gain amplifier, widely used in this application, has two channels that can accept a linear time sweep and produce an exponentially increasing gain with a 48-dB range (power ratio approaching 100,000:1).

The A/D Converter

There are many sources of noise that combine at the input of the ADC, including body tissue, gain stages and cable noise. As the last link in the chain, it is important that the ADC itself have low noise. Its noise must not be confused with the surviving signal coming from the other components. Quantization noise is improved by using higher resolution converters. Many ultrasound systems use 10-bit converters with theoretical quantization noise of -61.7dB. Newer machines are using 12-bit converters, which bring the theoretical quantization noise down to <-73dB.

Many ultrasound designers are concerned with harmonic distortion and artifacts at frequencies close to the fundamental. Unlike a state trooper's Doppler radar, which deals with a large frequency shift when measuring the velocity of a speeding Honda, the Doppler modes of an ultrasound system measuring the velocity of blood in a vein or artery produce a shift of only a few hertz. In the FFT plot, the areas near the base of the fundamental frequency spike must be very quiet and free of spurious signals, often caused by ADC or system clock jitter, so as not to mask out this shift. Linearity of the converter is also important to the quality of Doppler ultrasound.

Low intermodulation distortion in an ADC will help to prevent the various harmonic artifacts of Doppler return from mixing to form aliases or adding to form large spurs. The reflected signals inside the body can be considered as multi-tone signals. If the ADC has poor harmonic distortion characteristics, the tones will combine with the ADC's harmonics, which could dwarf a low-amplitude return signal.

Many ultrasound manufactures use 4× oversampling for improving signal-to-noise and to reduce complexity in antialias filters. However, a 12-MHz mammography machine will need better than 48 MHz to accommodate the system. Oversampling rates are dictated by the ability of the signal processing chain to process the data stream.

Display

Once the points have been scanned, they must be displayed. Now consider how the machine places the images on the screen. It will calculate the location of a target on the screen based on the time delays from element to element in the row of ceramics in the hand piece. It judges depth based on how long it took the signal to come back from each ceramic element. The pixel values will be read out of memory and modulate the CRT trace.

The machine must compute the location of each point and add color. Perhaps it will average several received scans together. Then it will start the CRT sweep at the top of the fan-shaped display.

Harmonic Imaging

To gain the improved resolution from higher-frequency, while ameliorating the dilemma of depth of penetration vs. energy level, harmonic imaging is used. Harmonic imaging gathers increased resolution by processing the second harmonic of the fundamental transmit pulse. The harmonic is generated by the tissue itself or the use of contrasting agents injected into the tissue. This technology will put the pressure on amplifiers and ADCs to minimize additional harmonics by maintaining low harmonic distortion.

Future Component Requirements

There is an ongoing demand for lower power components. In earlier days in hospitals, portable meant that the bulky machine had big wheels, and that it could be powered from the 120-V/15-amp receptacle in a hospital room rather than the 220-V/30-A receptacles in Radiology. Nowadays, there is growing interest in installing ultrasound in emergency vehicles and making it really portable. Tendencies in component design support such momentum. For example previous high-speed 10-bit ADCs drew > 400 mW. That's a lot of power when there are 256 converters in

close quarters. In contrast, the 10-bit, 40-MSPS [AD9203](http://products.analog.com/products/info.asp?product=AD9203) <<http://products.analog.com/products/info.asp?product=AD9203>> draws just 75 mW.

Cost has come down more than 2 to 3× from levels of a few years ago. This makes practical the use of higher-resolution, faster ADCs, such as the low-cost, 12-bit 65-MHz [AD9226](http://products.analog.com/products/info.asp?product=AD9226) <<http://products.analog.com/products/info.asp?product=AD9226>>.

More Things to Come

As time goes on it is logical to expect better images for less money. This will be made possible by ADC's with even higher resolutions and data rates. Many more samples can be made of reflected images as they arrive at the processor.

3-D imaging is now being developed. With these machines one can get a better overall view of an image, which can result in quicker and more accurate diagnosis and less unnecessary surgery.

More Information

Most ultrasound manufacturers maintain web sites <http://dir.yahoo.com/Business_and_Economy/Companies/Health/Medical_Equipment/Ultrasound/>. There are many technical articles and application notes available <http://dir.yahoo.com/Health/Medicine/Medical_Imaging/Ultrasound/>. Logistics prohibited the reproduction of ultrasound images here, but there are many available in cyberspace¹. Be sure and check out 3-D images, they are awesome. So are the prospects for more widespread use of ultrasound as a rapidly maturing noninvasive imaging technology. ▣

¹<http://home.hkstar.com/~joewoo/joewoo2.html>

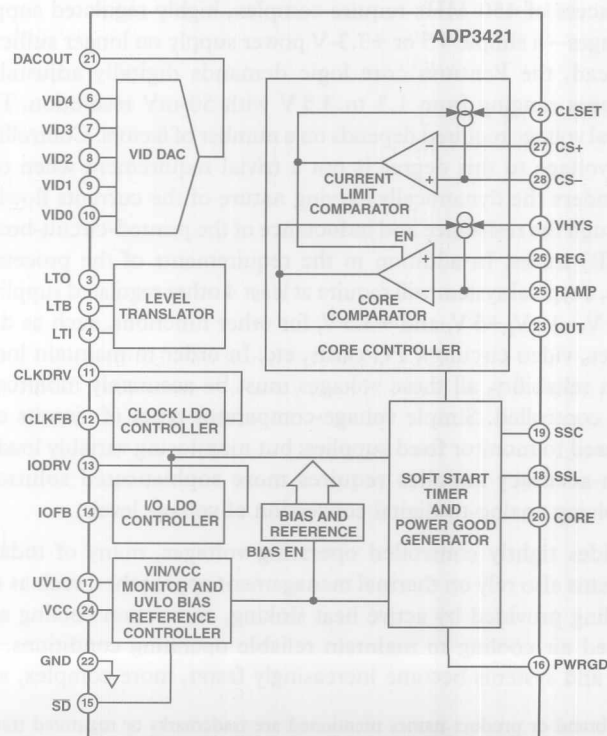
The author is indebted to Eberhard Brunner for conceiving the drawings used here.

New-Product Brief

POWER MANAGEMENT

Geyserville-Enabled DC/DC Converter Controller for Mobile CPUs

The *ADP3421* is a hysteretic voltage-programmed dc-to-dc buck converter controller with two auxiliary linear regulator controllers. It provides a total power conversion-control solution for a microprocessor, by delivering the core, I/O, and clock voltages. The main output voltage is set by a 5-bit VID code. The *ADP3421* features high-speed operation to minimize the size of the external inductors, resulting in the fastest change of current to the output. It also features optimally compensated active voltage positioning to minimize the size of the output capacitors while providing superior load transient response. The *ADP3421* is specified for a 3.3-V supply, and it draws only 10 μ A max in shutdown. The main output signal interfaces with the *ADP3410* dual MOSFET driver. The *ADP3421* is packaged in a 28-lead TSSOP, and operates over the 0 to +100° C temperature range. Price (1000s) is \$2.50.



A Chip You Can Use to Monitor Environmental Conditions on PC Motherboard Designs

by Matt Smith

Management of the electrical, mechanical, and thermal environment is of growing importance in today's microprocessor based systems. This article will focus on accurate voltage and temperature monitoring as well as serving other hardware monitoring requirements. We will consider techniques for making accurate measurements and offer solutions for the range of hardware monitoring tasks in Pentium II-based systems as embodied in the ADM9240 chip.

As designers of ICs and systems seek to squeeze every last morsel of performance out of their designs, hardware monitoring and control have become an integral part of circuit board design objectives. Examples include the need to maintain accurate supply voltage levels and continually dispose of the heat generated by high-performance chips. The feedback portion of the control loop is accomplished by *hardware monitoring*—the continuous measurement of critical system parameters, such as power supply voltages, internal temperatures, cooling fan performance, and other environmental factors. System performance can be maximized by closely controlling these parameters to remain within tight limits so as to maintain optimum operating conditions for the circuitry and avoid reduction of component life.

The latest generation of Intel¹ Pentium-based products clearly demonstrates the importance of hardware monitoring and control. The newest Pentium II microprocessors running with clock rates in excess of 450 MHz require complex, highly regulated supply voltages—a simple +5 or +3.3-V power supply no longer suffices. Instead, the Pentium core logic demands digitally adjustable voltages ranging from 1.3 to 3.5 V with 50-mV resolution. The actual voltage required depends on a number of factors. Controlling the voltage to this degree is not a trivial requirement when one considers the dynamically varying nature of the currents flowing through the resistance and inductance of the printed-circuit-board (PCB) traces. In addition to the requirements of the processor chip, a typical system will require at least 4 other regulated supplies, +12 V, -12 V, +5 V, and +3.3 V, for other functions, such as disk drives, video circuitry, PC cards, etc. In order to maintain long-term reliability, all these voltages must be accurately monitored and controlled. Simple voltage-comparator types of circuits can be used to monitor fixed supplies; but monitoring variably loaded high-accuracy supplies requires more sophisticated solutions involving analog-to-digital conversion of voltage levels.

Besides tightly controlled operating voltages, many of today's systems also rely on thermal management approaches, such as the cooling provided by active heat sinking, convection cooling and forced air cooling to maintain reliable operating conditions. As ICs and systems become increasingly faster, more complex, and

more dense, removing the excess heat and maintaining safe, reliable operating temperatures has become increasingly important. Temperature sensing, often coupled with fan-speed monitoring and control, are a couple of the techniques being employed today to ensure system reliability. By controlling fan speed, greater efficiency, reduced power dissipation and lowered noise levels are achieved.

Another important area that benefits by effective hardware monitoring is total cost of ownership (TCO). All the vital functions are monitored continuously and the results communicated to the systems management software. Impending failures can be detected, the sources identified, and corrective action taken—or even system shutdown invoked—before expensive damage occurs. For example, a clogged-up cooling fan may be detected by monitoring its speed. When the speed has decreased by 10 to 15% from its nominal speed, the software can note the problem and shut down the system before the deterioration has caused additional damage. Replacing a fan for \$10 is more appealing than replacing a \$1000 CPU or an even more-expensive system board.

Multi-channel voltage, temperature and fan-speed monitoring, together with programmable limit-setting for each of these parameters, goes a long way towards meeting the monitoring and control objective. We will discuss below some specific techniques for achieving this in Pentium-based systems. The demonstration will employ a new monitoring IC, ADM9240, from Analog Devices, to show approaches to optimizing the monitoring strategy.

Voltage Monitoring

As many as six voltage channels require monitoring in a typical system. Typical system supplies include a combination of some or all of the following: +12, -12, +5, +3.3, +2.7, and +2.5 V.

With so many (and different) supplies to keep track of, a multiplexed data-acquisition system with digital readout provides the greatest flexibility. An A/D converter-based solution facilitates software control and limit setting. Once converted to the digital domain, the data is easily manipulated, processed, and stored for historical reference.

Here are a few of the considerations that must be addressed so that the signals are accurately converted from analog to digital: Since the supplies being measured are usually generated using switched mode techniques, noise introduced by the switching can make their voltage difficult to monitor accurately. Switching glitches and load-dependent voltage excursions can be a source of spurious alarms. Thus it is important that the monitoring circuitry reject supply glitches and excursions but still be fast enough to detect when the supply is really out of tolerance. When the supply is indeed out of tolerance, it is important to report it so that the situation can be dealt with as quickly as possible to avoid errors in system performance or even damage.

The input circuitry of the ADM9240 (Figure 1) serves the dual role of

- (a) Filtering the input signals.
- (b) Attenuating the input levels to scale them to the reference voltage of the on-board ADC.

Having the attenuation network integrated on-chip provides an important advantage. Any errors it introduces due to inaccurate resistors or mismatch are already included in the specifications for the channel, so the user does not need to further increase the system error budget.

¹All brand or product names mentioned are trademarks or registered trademarks of their respective holders.

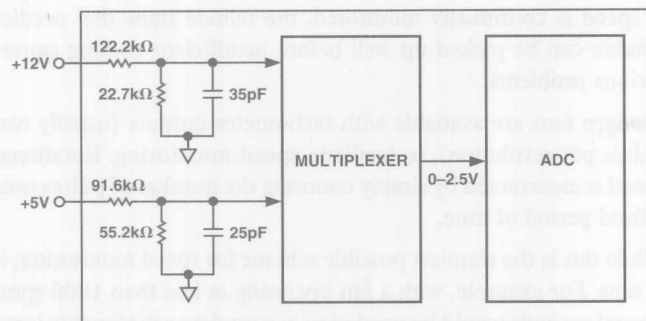


Figure 1. Input channel filtering & attenuation.

The input range on the ADM9240 is also biased so that nominal input voltage levels correspond to 3/4 full-scale on the ADC (Figure 2). This scaling provides a range of from +25% overvoltage to total failure. Having most of the dynamic range at the lower end takes into account the majority of cases of error and also allows greater flexibility, since it is possible to monitor lower voltage supplies (specified at less than the standard levels listed), but with some reduction of accuracy.




Figure 2. Input transfer function.

Monitoring the Core Voltage

Besides monitoring the fixed supplies, Pentium II-based systems also require accurate monitoring of the processor core voltage, VCCP. Today's Pentium IIs (P2) use a 5-bit VID (voltage identification) code (up from 4 bits on previous-generation products) Depending on the VID code provided by the P2, the core voltage can be set anywhere between 1.3 V and 3.5 V.

VID Code Table

VID4	VID3	VID2	VID1	VID0	Voltage (V)
0	1	1	1	1	1.30
0	1	1	1	0	1.35
0	1	1	0	1	1.40
⋮	⋮	⋮	⋮	⋮	⋮
1	1	1	1	1	No CPU
1	0	0	1	0	3.3
1	0	0	0	1	3.4
1	0	0	0	0	3.5

The voltage monitoring requirements discussed earlier apply also to the VCCP supply, but the tolerances are much tighter. The A/D converter input range for monitoring this is set at 0 V to 3.6 V with 3/4 full scale at 2.7 V. This provides sufficient dynamic range and accuracy to accommodate other processor core voltages, even beyond P2 requirements.

In dual processor systems, which may employ different processor core voltages, the ADM9240 makes available a second multiplexed input channel (VCCP2).

Monitoring Negative Voltages

Negative voltages can be monitored on positive input channels by inverting the signal's polarity. But this may not be cost effective—it requires an inverting op amp and wastes chip "real estate." A lower cost scheme, using positive bias and an inverted interpretation of the range, can also be utilized. This is illustrated in Figure 3: Resistor R2 is biased up to +5 V; and the upper and lower limits (as well as the overvoltage (75%) and undervoltage (25%) ranges) will be transposed. Since the offset voltage will be dependent on the +5-V reference level, either an accurate +5-V reference should be used or—if the 5-V supply itself is used—this input should be measured first and the -12-V supply's limits set accordingly.




Figure 3. Monitoring a -12-V supply using a positive input channel (VCCP).

Temperature Monitoring

Temperature monitoring enhances reliability as well as allowing the efficiency of a close approach to maximum performance. It can also serve to protect the system against overheating if the cooling system fails completely or deteriorates to the point of inadequacy.

Silicon sensors are becoming increasingly important as temperature transducers in electronic systems because they are linear, accurate, cheap, reliable; and can be incorporated on the same IC as other analog or digital functions. They take advantage of the relationship between base-emitter voltage (V_{BE}) and current density (current/emitter area) in silicon bipolar junction transistors to generate a voltage proportional to absolute temperature (PTAT). If currents in a fixed ratio, r , flow through two identical transistors (or if equal currents flow through one transistor and a set of r identical paralleled transistors) the differential V_{BE} is PTAT. Figure 4 is a circuit that illustrates the principle. A_1 and A_2 are emitter areas, I_S is reverse saturation current, and k/q is the ratio of Boltzmann's constant to electron charge, about 86 $\mu\text{V/K}$.

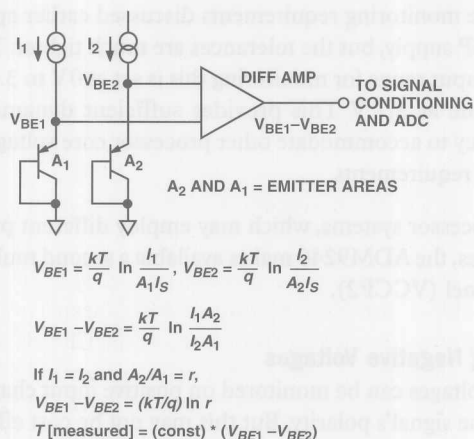


Figure 4. How a silicon temperature sensor works.

On the ADM9240, on-chip temperature sensing uses an additional multiplexer channel (Figure 5). When cycling, the analog inputs and the temperature channel are each selected in turn by the multiplexer and converted into a digital quantity by the ADC.




Figure 5. Voltage & temperature monitoring.

Offsetting, scaling and data manipulation provide a two's-complement output. Although the theoretical temperature span is from minus 128°C to +127°C, practical device and package constraints limit it to about -40°C to +125°C.

The location of the temperature sensor is important to accurate temperature measurement. Ideally it should be in intimate physical contact with the object being measured. This is not always possible, especially when a single sensor is just one function on a multifunction IC where other considerations must be taken into account. If direct thermal contact is not possible, it is important to characterize the difference between the temperatures of the sensor and the desired measurement point. In this way, a known offset may be used to compensate for the temperature difference.

Fan Speed Measurement

Fan speed sensing provides an invaluable early warning signal of potential problems. Fans are a weak mechanical component in an otherwise highly reliable electronic system. While modern brushless fans are much more reliable than earlier brush types, they are still prone to mechanical wear and tear. Bearing wear and increased friction slow the fan's rotational speed, resulting in reduced air.

If speed is continually monitored, the telltale signs that predict trouble can be picked up well before insufficient cooling causes serious problems.

Modern fans are available with tachometer outputs (usually two pulses per revolution), to facilitate speed monitoring. Rotational speed is ascertained by simply counting the number of pulses over a fixed period of time.

While this is the simplest possible scheme for speed monitoring, it is slow. For example, with a fan operating at less than 1000 rpm, several seconds would be needed to accumulate a reasonably large and accurate count.

The technique employed on the ADM9240 does not count the fan tachometer output pulses directly. Instead it uses the tachometer output as a gating signal for a high-frequency internal clock. By counting the number of gated pulses, the fan's period may be determined. The accumulated count is proportional to the fan's tachometer period and inversely proportional to speed.

Specifically, an on-chip 22.5-kHz oscillator is gated into the input of an 8-bit counter for two periods of the fan tachometer output, corresponding to the time for one revolution of the fan (Figure 6).

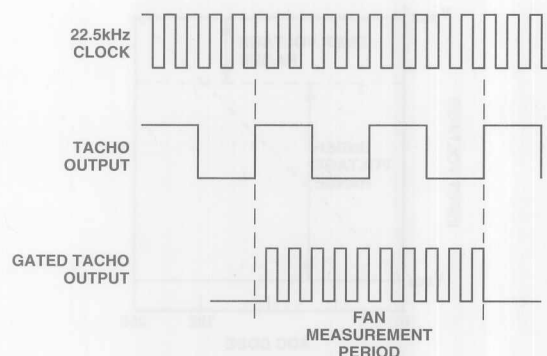


Figure 6. Tachometer output & speed determination.

To accommodate fans of different speed, a pre-scaler (divisor) may be added before the counter. Consider the following example using the ADM9240.

With a divisor of 2, a fan with two output pulses per revolution running at 4000 rpm, gives a count of 168. That is, at 4000 rpm, there are 8000 pulses per minute, 133.3 pulses per second; so the interval between pairs of pulses is 15 ms, giving a count of $0.015(22,500/2) = 168+$.

As the fan slows down, the count increases to the counter's maximum count of 255, which occurs at

$$4000 (168.75/255) = 2647 \text{ rpm.}$$

Interfacing Fan Tachometer Outputs to 5 V/3 V Logic

Because fans are generally powered from voltages higher than the logic/monitoring circuitry it is necessary to provide an interface that does not overstress the logic or forward bias some internal junctions. Voltage clamping using a resistor/Zener diode network provides a good solution (Figure 7). The Zener breakdown voltage should be chosen so that it is lower than the power supply voltage to the logic. With 5-V logic, a 4.0-V Zener is suitable.




Figure 7. Tachometer output interface.

Controlling the Fan Speed

If the fan moves more air at rated voltage and ambient temperature than is needed for adequate cooling, its speed may be controlled to reduce acoustic fan noise and power consumption while maintaining the temperature at a safe level.

The simplest form of control is linear adjustment of the supply voltage to the fan. For example the speed of a 12-V fan may be limited by adjusting the supply to voltages less than 12 V.

However, one must take into consideration that the fan may not start up reliably if the supply voltage is fixed at a low value. With a D/A converter to vary the speed (Figure 8), the fan can be started at a higher speed, then slowed down to the correct value. With a 12-V fan, the minimum reliable operating voltage may be as high as 6 or 7 V, allowing for a considerable range of adjustment.

The ADM9240's 8-bit DAC can be used for fan speed control. The 1.25-V output of the DAC will need an external amplification/current-boosting stage to drive the fan.




Figure 8. Fan speed control.

Control and Monitoring Over a LAN Network

Hardware monitoring and control may be further extended to operate over a network so that the health of an entire network of computers can be continually monitored. Intel's LAN Desk Client Manager (LDCM) is an example of network management software which can monitor and troubleshoot individual workstations over a network, as well as providing early warning to the systems administrator of potential future problems.

Many network problems occur as consequences of the installation of incompatible software—or hardware—by an inexperienced user. Both can be monitored remotely over the LAN. Chassis intrusion sensors can be used to detect unauthorized tampering with a system. Typical sensors include simple microswitches, reed switches, Hall-effect switches or even optical sensors. When the case is opened, the switch is toggled or the optical beam is broken.

The ADM9240 includes an input line, which can be connected to a chassis intrusion switch to alert the monitoring system. Generally the switch is wired to a latching circuit, using a flipflop or a thyristor. For this reason, the chassis intrusion latch must be resettable by the systems administrator. On the ADM9240, the chassis intrusion line can also be temporarily configured as an output line so that a Clear pulse may be sent to clear the latch.

Figure 9 shows a block diagram of the complete ADM9240. With its combination of voltage, temperature, fan, and chassis intrusion monitoring, a better-controlled operating environment is available for the electronics. The benefits to the user are increased stability, reliability and reduced ownership costs.

Figure 10 illustrates a complete monitoring solution using the ADM9240. This circuit is suitable for Pentium II-type motherboards. All six power supplies are simultaneously monitored for either overvoltage or undervoltage conditions which would pose a threat to the electronics. The high and low limits are programmed over a 2-wire Systems Management Bus (SMBus). The master controller is generally a PIIX4 Southbridge chip but could also be a dedicated microcontroller. In addition to voltage monitoring, the circuit monitors the speed of a pair of cooling fans via J2 and J3. One of these fans (J3) is being speed controlled to limit acoustic noise using the DAC on the ADM9240, while the second fan runs continuously at full speed. Linear speed control provides a reliable,

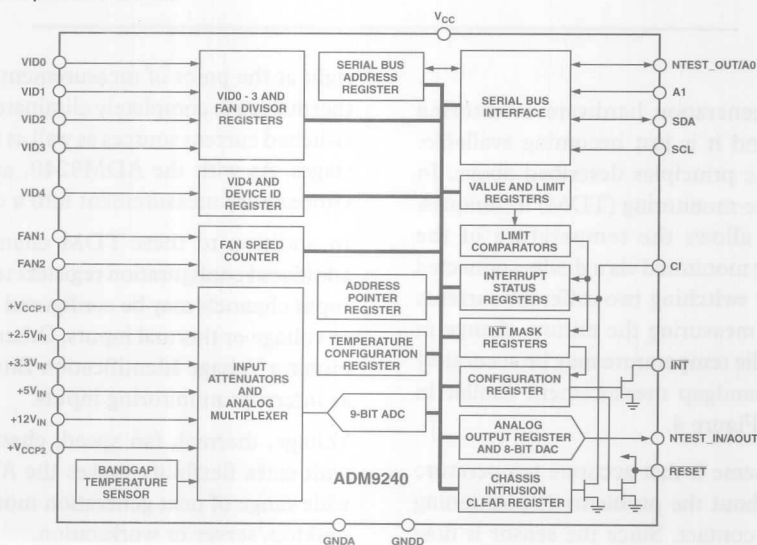


Figure 9. ADM9240 block diagram.

low-noise solution to maintaining low acoustic emissions. Unauthorized tampering with the system is detected and latched via an optical, mechanical or magnetic switch—appropriately located to avoid tampering—connected to J1. Note that the

detection and latching circuitry is being powered from a backup battery so that monitoring continues even when the system is unplugged. When asserted, the latching logic may be cleared by the systems administrator, using a special command to the ADM9240.

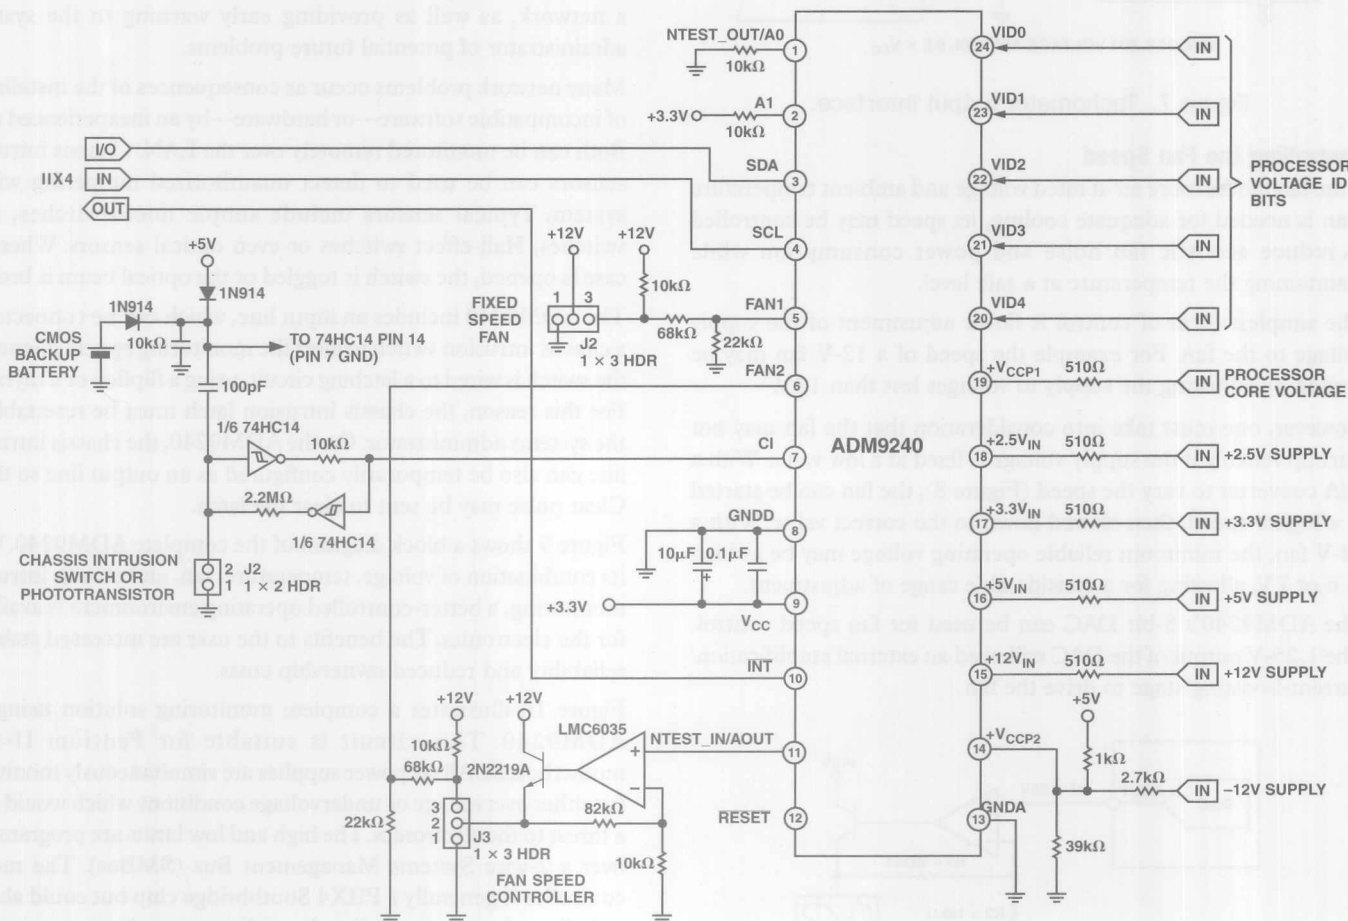


Figure 10. Complete monitoring solution using ADM9240.

The Next Generation

As this was written, the next generation hardware monitoring solution was in development; and it is just becoming available. The *ADM1024* embodies all the principles described above. In addition, it includes thermal diode-monitoring (TDM) techniques. This revolutionary technique allows the temperature of the Pentium die itself to be constantly monitored via a diode-connected transistor on the Pentium II. By switching two different currents through the on-chip diode and measuring the minute change in the diode's forward voltage, the die temperature may be accurately ascertained by a multiplexed bandgap measurement similar in philosophy to that described in Figure 4.

The principal benefit of this scheme is that accurate temperature measurements are obtained without the problems of positioning an external sensor for intimate contact. Since the sensor is now

right at the point of measurement, it can be highly accurate, and thermal lag is completely eliminated. The *ADM1024* contains the switched current sources as well as filtering and amplification input stages. As with the *ADM9240*, an on board ADC converts the temperature measurement into a digital reading.

In addition to these TDM channels, the *ADM1024* contains additional configuration registers to provide even greater flexibility. Input channels may be configured as needed to measure fan speed or voltage or thermal inputs. Other channels may be configured to monitor Voltage Identification Bits (VID) or indeed may be used as interrupt monitoring inputs.

Voltage, thermal, fan speed, chassis, VID monitoring combined with extra flexibility makes the *ADM1024* suitable for use on a wide range of next generation motherboard designs whether it be desktop, server or workstation. ■

Measuring Temperatures on Computer Chips with Speed and Accuracy

A new approach using silicon sensors and off-chip processing

by Matt Smith, Staff Engineer, Interface Products, ADI Limerick, Ireland

Silicon Temperature Sensors

Silicon sensors are becoming increasingly important transducers in electronic systems. As systems become more complex, more compact, and denser—and run faster and hotter—it becomes increasingly vital to monitor critical temperatures. Traditional sensor techniques, such as thermocouples, thermistors, and RTDs, are now being displaced by silicon sensors, with their ease of integration and use. Many traditional sensor types are inherently nonlinear and require signal conditioning (i.e., compensation, look up tables, excitation circuitry, etc.) to accurately convert temperature into an electrically measurable quantity such as voltage or current.*

Silicon sensors, on the other hand, are linear, accurate, low-cost, and can be integrated on the same IC as amplifiers and any other required processing functions. The actual sensing element in a silicon sensor is a simple P-N transistor junction. The voltage across a regular P-N transistor junction has an inherent temperature dependency of about $2 \text{ mV}/^\circ\text{C}$ and this fact may be used to develop a temperature measuring system. Silicon sensors are new by sensor-industry standards but are very mature by semiconductor-industry standards. For example, the **AD590** $1\text{-}\mu\text{A}/^\circ\text{C}$ IC sensor was introduced more than 20 years ago!†

In order to separate the variation with temperature from the effect of current level and remove offsets, the most common technique is to base the measurement on two transistor junctions. By operating two identical transistors at a constant ratio of collector current densities, r , the difference in their base-emitter voltages will be $(kT/q)(\ln r)$. Since both k (Boltzmann's constant) and q (electron charge) are physical constants, the resulting voltage is directly proportional to absolute temperature, T (PTAT).

Temperature-monitoring products available from ADI, which incorporate this type of temperature sensing, usually integrate it with additional functionality. For example, it may be combined with analog-to-digital conversion circuitry. Figure 1 shows a block diagram of the **AD7415**; it contains temperature-sensing circuitry,

an amplifier, and an ADC, along with a two-wire I²C interface. Other products, such as the **ADM9240**, which was featured in *Analog Dialogue 33-1*, include many additional functions, such as voltage monitoring and fan-speed monitoring, as well as on-chip limit setting.

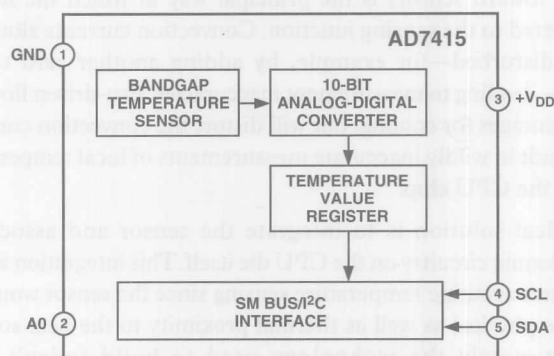


Figure 1. Temperature sensor plus ADC.

Sensor Mounting Considerations—THE PROBLEM

While a silicon sensor is a very accurate temperature transducer, it is important to remember that it will only measure its own junction temperature, and thus its own *die* temperature. This is fine if one is simply interested in monitoring approximate zone temperatures within an enclosure or environmental temperature (and convection and conduction conditions are adequate). If, however, one must monitor the local temperature within a heat source or a computer chip, such as a Pentium® III CPU, or a high performance graphics chip, much more is at stake and the situation is not quite so straightforward. In order to get an accurate measure of the temperature of the heat source, the sensor must be in close proximity to the source itself. The accumulation of thermal resistances between the sensor and heat source will lead to measurement errors and uncertainties. The physical mounting problems required to be solved in order to get accurate temperature measurement may be simply impossible to deal with in many situations, resulting in derating and suboptimal performance.

For example, if IC temperature sensors had to be mounted on the circuit board, it is very unlikely that they could be in close physical contact with the “hot spot” of the object being monitored. It might be possible to work around the mounting difficulties with tiny two and three terminal devices, but with multi-lead packages it is virtually impossible.

Offset Calibration?

One approach might be to add a well-chosen offset to account for the temperature difference between the sensor and the heat source. The required offset can be derived during system characterization by comparing the displayed temperature with the actual temperature. Since the offset needed at room temperature will almost certainly be different from the offset required at elevated temperatures, a simple offset register is generally not enough. A lookup-table approach is one way of working around the problem. This approach might be effective, albeit unwieldy, for a fixed system, but the look up tables would differ when the system configuration changes.

*Practical examples of signal-conditioning considerations, designs, and circuits can be seen in the seminar notes, *Practical Design Techniques for Power and Thermal Management*, Section 6. For information about this book and how you can purchase it, phone (617) 461-3392 or e-mail lit.center@analog.com.

†*Analog Dialogue 12-1*, 1978, pp. 3-5. See also M. P. Timko, “A two-terminal IC temperature transducer,” *IEEE Journal of Solid-State Circuits*, vol. SC-11, 1976, pp. 784-788.

Pentium is a registered trademark of Intel Corporation.

For example, consider trying to measure the temperature of a CPU on a motherboard by placing a temperature sensor as closely as possible to the CPU. The sensor will most likely be at least 1 cm away from the heat source (CPU). The thermal resistance of the path through the board material between the two is very high, and air currents (i.e., convection or fan-driven flow if directed from source toward sensor) is the principal way in which the heat is transferred to the sensing junction. Convection currents alone are easily disturbed—for example, by adding another card to the system—leading to measurement inaccuracies. Fan-driven flow has its advantages for cooling, but will distort the convection currents and result in wildly inaccurate measurements of local temperature within the CPU chip.

The ideal solution is to integrate the sensor and associated conditioning circuitry on the CPU die itself. This integration would guarantee accurate temperature sensing since the sensor would be in close physical as well as thermal proximity to the heat source. Unfortunately the technology used to build today's high performance CPUs is not compatible with the technology used to build highly accurate temperature sensors and associated amplification circuitry.

The Answer: Sense the CPU Directly

The best approach to the problem is to provide P-N-junction sensing on the CPU die near the hot spot(s)—and then use an external conditioning IC to do the rest. This approach allows CPU temperatures to be directly measured without any uncertainty. The newest Intel Pentium II and Pentium III CPU's contain an on-chip thermal diode monitor (TDM) to facilitate this. On Slot 1 CPUs, two pins, THERMDP and THERMDN, provide access to the on-chip diode. To supply signal conditioning and convert the minute voltage changes into robust measurable results in digital form, a new generation of products from Analog Devices, the ADM102x series, supply the required conditioning and conversion circuitry.

TDM to Digital—A New Approach

The trick is now to translate the minute voltage changes due to temperature into really measurable signals and digitize them. The low signal levels would by themselves pose a difficult instrumentation problem, but it is further complicated by the noisy environment that the circuit must operate in. Picture if you will the electrical environment within a digital computer chip! The signal could very easily be swamped by the noise making it impossible to recover the signal. Also, manufacturing variations from unit to unit cause differences in junction response. We will now discuss how the technique works, how it compares with more traditional techniques and how to extract optimum performance from it.

THE SOLUTION

First, for a given current level, the absolute forward voltage drop of the diode isn't very well controlled in the CPU manufacturing process. Also, because voltage depends on *absolute* (i.e., Kelvin) temperature, the forward voltage value is many times larger than the change in its value per 1°C temperature change. Therefore, the most important requirement is to remove the absolute value of the diode voltage from the equation before any amplification can occur.

Individual device calibration is an option but not a practical one. Rather, a technique comparable to the two-transistor approach described above is used, except that the ratio of current density (current per unit area), r , depends on changing the current in the same diode instead of using the differing areas of two diodes with equal currents. This technique, called "delta- V_{BE} calibration," forces two different levels of current through the thermal diode junction and measures the change in forward voltage. The first current may be considered as a calibration current and the V_{BE} forward voltage value of the junction is ascertained. The V_{BE} value is then measured again with a second current. The change or difference in V_{BE} is proportional to absolute temperature. It is independent of the junction's forward voltage or other differences due to manufacturing variations.




Figure 2. TDM monitoring.

$$V_{BE} = (kT/q) \ln(I_C/I_S)$$

Since I_S is a property of the transistor and is unchanged for either current,

$$V_{BE1} - V_{BE2} = \Delta V_{BE} = (kT/q) \ln(I/NI) = (kT/q) \ln(1/N)$$

Since N , k , and q are all known constants,

$$T = (\text{Constant})(\Delta V_{BE})$$

The output from the ΔV_{BE} sensor varies at approximately 2.2 mV/°C. This signal requires conditioning and amplification. The actual ΔV_{BE} sensor is shown as a substrate transistor since this would be the case in practice for an on-chip junction. It could equally well be a discrete transistor. If a discrete transistor is used, the collector will not be grounded and should be linked to the base. To help prevent ground noise from interfering with the measurement, the more negative terminal of the sensor is not referenced to ground, but is biased above ground by an internal diode at the D- input. To measure ΔV_{BE} , the sensor is switched between operating currents of I and $N \times I$.

Filtering and Amplifying

The resulting waveform is passed through a 65-kHz low-pass filter to remove noise, then to a chopper-stabilized amplifier that performs the functions of amplification and rectification of the waveform to produce a dc voltage proportional to ΔV_{BE} . This voltage is measured by the ADC to give a temperature output in 8-bit twos-complement format. To further reduce the effects of noise, 16 measurements are made, the results are averaged, and the average result is then provided at the output.




Figure 3. Signal conditioning.

So how good is the TDM approach in practice?

It is interesting to compare a TDM measurement with a more traditional thermistor approach. The following example compares results obtained using a thermistor and a TDM channel to measure the temperature of a 333-MHz Pentium II in a Slot 1 cartridge. The thermistor is in direct physical contact with the cartridge heatsink. The TDM channel uses the on-chip diode and an ADM1021 (with circuitry similar to those discussed above) to provide the signal conditioning.

Besides being more accurate, it does not suffer from thermal lag. While contact with the heat sink is superior to earlier approaches where contact was not even made, it still exhibits major disadvantages. As may be seen in Figure 4c the thermistor lag completely misses many of the thermal events due to its slow response time. Figure 4a shows a power up event while Figure 4b shows a power-down. Errors in excess of 20°C (representing cartridge temperature instead of actual chip temperature) are evident.

Even more significant is Figure 4c where the CPU is cycled in and out of Suspend mode. The thermistor completely misses these 20°C thermal events. It's easy to see how it would fail to protect a system in the event of a rapid temperature rise due to a fault condition. All plots also demonstrate the offset error (due to package temperature drop) between the TDM and the thermistor as the temperature increases. The offset can be taken care of by system calibration but there is nothing one can do to compensate for thermal lag. Indeed if additional system cooling were employed, the errors between the TDM and the thermistor would be greater still.




Figure 4a. Start up.




Figure 4b. Shut down.




Figure 4c. Suspend/wake up cycling.

Using Discrete Transistors for TDM

So the TDM approach is very effective if the sensing diode is integrated onto the die of the CPU whose temperature is being measured. What about using this approach to measure temperatures where there isn't an on-chip TDM, or to measure the temperature of heat sources other than IC's? The ΔV_{BE} TDM approach may also be used with stand-alone discrete transistors. Any NPN or PNP general-purpose transistor, such as 2N3904 or

2N3906, may be used as a remote sensor. With a discrete transistor, connect the base to the collector to form a 2-terminal device. Transistors are good temperature sensors since they have low thermal mass and are easily mounted.




Figure 5. Transistor TDM.

If the transistor sensing junction is a significant distance away (>6 feet) and if it is used in noisy environments, the best method to preserve signal integrity and prevent interference is to use twisted shielded cable. The maximum cable length is limited by cable capacitance and by series resistance. Capacitance between D+ and D- causes settling time errors since the switched current needs to have fully settled before the conversion is made.

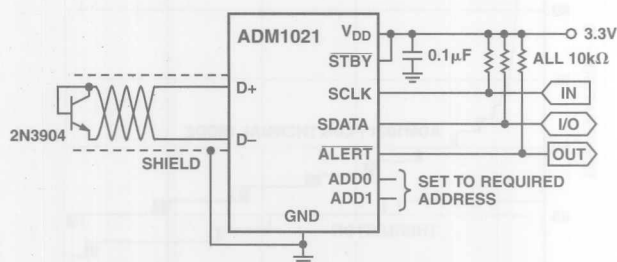


Figure 6. Transistor TDM plus ADM1021 for remote sensing.

TDM in Noisy Environments

It is very important to observe some guidelines when utilizing thermal diode sensing techniques, especially in noisy environments. The PC environment is inherently noisy and appears to be getting noisier as PCs get faster. As CPU speeds hurl towards 1 GHz, EMC noise becomes more of a headache. High speed graphics ports (AGP), high speed random-access memory, and high speed disk access mean that there are many opportunities and paths for noise to couple into sensitive analog circuitry. TDM is a very sensitive approach. The circuitry that drives the thermal diode consists of high-impedance, low-level current sources. To prevent interference, the TDM lines should be kept as short as possible and shielded if there are high frequency noise sources in the vicinity.

Additional Features on the ADM1021

In addition to the TDM channel, the ADM1021 includes an on-chip transistor for local or environmental temperature monitoring. A programmable conversion rate (from 1 conversion per 16 seconds up to 8 conversions per second) facilitates high update rates, where rapid temperature changes must be recorded. If fast updates are not required, then lower update rates may be used to conserve power.

The ADM1021 also contains four limit registers to store local and remote, high and low temperature limits. A functional block diagram is shown in Figure 7. A typical system configuration using the ADM1021 is illustrated in Figure 8.

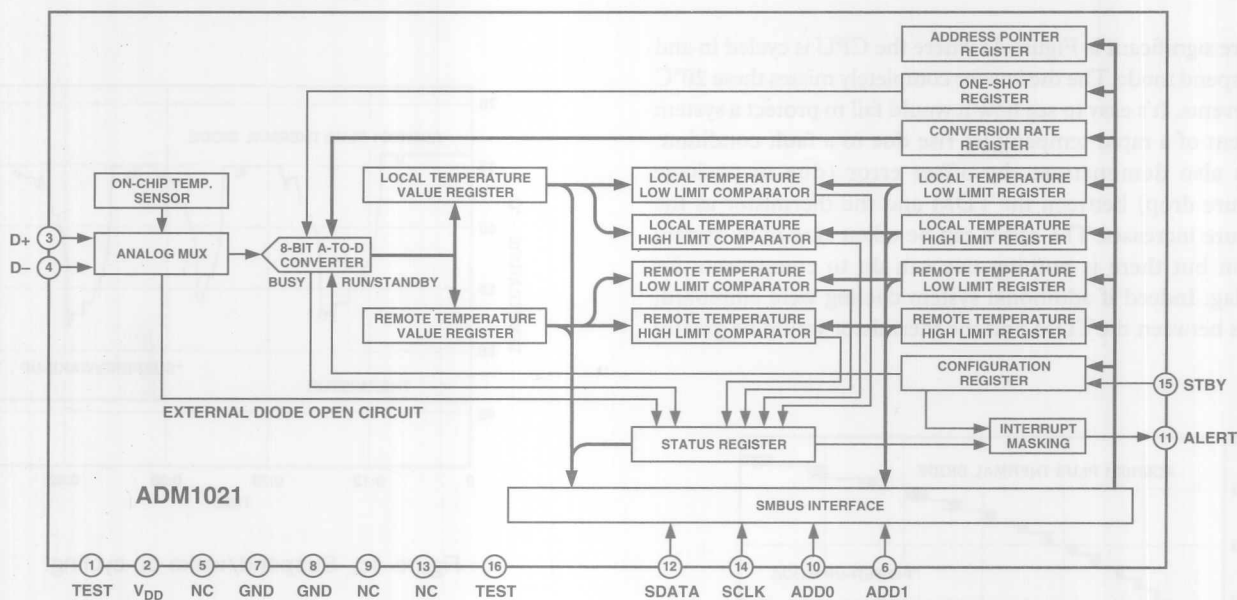


Figure 7. ADM1021 functional block diagram.

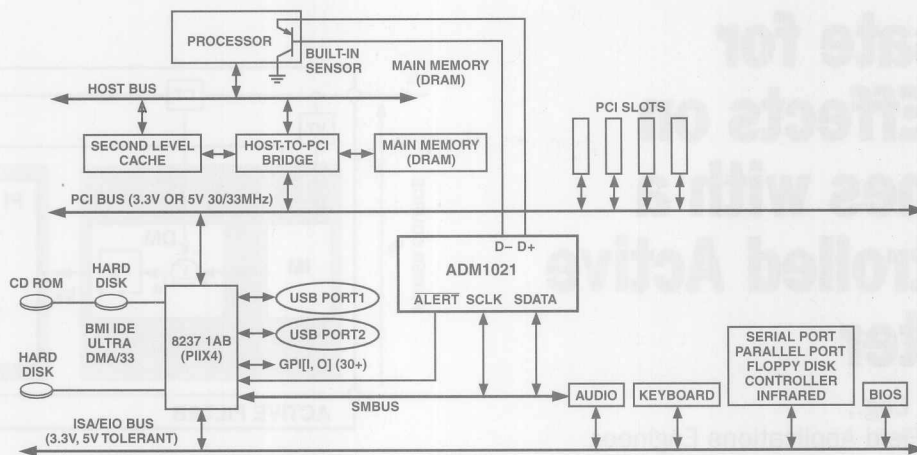


Figure 8. System architecture.

TDM CPU Monitoring Facilitates Optimum Cooling

A thermal profile of a real-life notebook computer is illustrated in Figure 9. This shows how the temperature ramps up on the chip and in the computer environment once the monitoring utility starts running (after power-on and Windows boot). It is interesting to note how hot both the CPU and the internal environment are running. The BIOS sets a CPU temperature limit of 92°C. When this temperature is reached, the fan is switched on and remains on until the temperature drops below 82°C. Because both high and low limits are programmed into the ADM1021, the fan can control the CPU temperature within a band between 82°C and 92°C. The temperature will oscillate between these two levels. If the fan fails, a higher temperature limit will shut the system down if it is breached. It is also interesting to note that the environment temperature also reaches very high temperatures, approximately 10°C below the CPU temperature, within a notebook housing.

This example illustrates the importance of TDM techniques in CPU temperature management. Prior to this technique, it would have been impossible to extract such high performance levels from the CPU without overheating—or unduly wasting battery life by continuously cooling the system. □




Figure 9. Temperature profile from a typical notebook computer.

New-Product Brief

AUDIO (ANALOG & DIGITAL)

AC '97 SoundMAX® Audio Codec

The AD1881 audio codec is used to add audio capabilities to computer motherboards when directly connected to the Intel® 82440MX Mobile Chipset or the Desk Top 82801 I/O controller hub (ICH). This direct connection capability makes use of existing system resources, adding flexibility, improving sound quality (including Phat™ 3-D stereo enhancement), and lowering system parts costs by eliminating components. The AD1881 was developed and validated to the Intel Audio Codec '97 (AC '97) 2.1

Component Specification. Systems designated as SoundMAX®, an Analog Devices registered trademark, include an audio codec, amplifiers and driver software, and several specific features for notebook and commercial desktop PCs. Specific features include variable sample-rate capability (7 kHz to 48 kHz \pm 1 Hz), mixer modes, true line-level output, multibit sigma-delta architecture for SNRs > 90 dB, 16-bit stereo full-duplex, capabilities for Line, CD, Video, and Aux (stereo), plus speakerphone, PC Beep, and Mic (\times 2) (mono). The AD1881, housed in a 48-lead LQFP, is specified for 5V supplies (3.3-V digital) and 0 to 70°C temperature range. Price (1000s) is \$4.64.

Compensate for Loading Effects on Power Lines with a DSP-Controlled Active Shunt Filter

Michał Gwozdz, Ph.D. Eng.,
Analog Devices DSP Field Applications Engineer,
P.E.P. ALFINE
Ryszard Porada, Ph.D. Eng.
Poznań University of Technology

INTRODUCTION

Systems with passive LC elements, such as capacity compensators, resonant passive filters of higher harmonics, or filters with structures and parameters determined by optimization methods have been traditionally used to compensate for power factor and other loading effects imposed by electricity users on the power network. However, widespread use of loads that involve power-electronic systems can cause major distortion of voltage and (especially) current waveforms, and even cause substantial dc currents to flow in power transformer secondaries. For these types of loads, the kinds of compensation systems mentioned above often prove to be unsatisfactory. Nowadays, power-system engineers are more likely to consider using other types of compensators, especially active power filters or hybrid systems (power filters with passive L-C elements such as those described in References 2, 3, 6, 7, 8, 9, 11) to increase system efficiency.

Recent approaches to development of compensation methods aim to develop compensators of a type that would be able to realize dynamic compensation (in real time) and would also be more resistant to interference caused by the power network or electricity users. Their objectives include optimization of the loads as seen by power sources (power network). According to Fryze's suggestion [5] and subsequent developments [4, 10, 12, 13], to achieve such compensation it is necessary to eliminate a differential current (between a distorted load current and an ideal form of current (i.e., in-phase sine wave)) flowing through the power source. In concept, this can be done by generating and injecting a current equal to and in opposite phase to the differential current. In practice, obtaining such a source is difficult; what is really called for is an active system with parametric elements or controlled-current power sources.

STRUCTURE OF AN ACTIVE FILTER

In this article we consider a proposal to employ a power-electronic current source controlled through the use of digital-signal-processing computer technology to achieve an *active shunt filter* (alternative names are: *differential current compensation system* or *compensator*), to approach a realization of optimal compensation. The assumed aim is dynamic compensation of differential current, which is the difference between load current $i_L(t)$ and reference current $i_{REF}(t)$. The reference current is the optimal active current calculated with the method suggested in article [10]. Figure 1 shows the block diagram of the system.

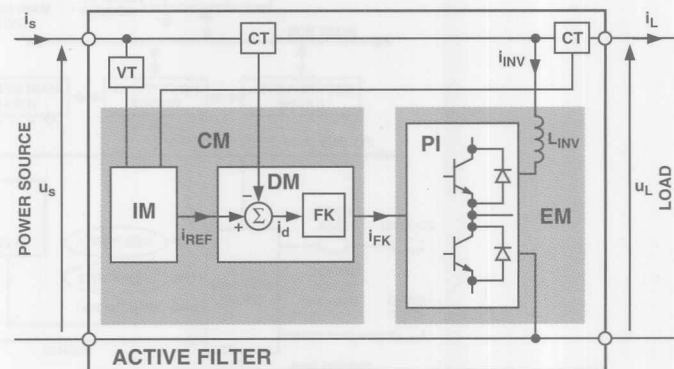


Figure 1. Block diagram of the active filter.

The active filter consists of the following modules and elements:

- Control Module (CM), based on a microcomputer system with digital signal processing (DSP),
- Execution Module (EM) in form of a power-electronic current source,
- Voltage (VT) and current (CT) transducers [Types LA55-P and LV25 (LEM®)].

The active filter control process occurs in two phases:

- Determining the reference current $i_{REF}(t)$,
- Dynamic shaping of desired compensator current in the form

$$i_{INV}(t) = i_L(t) - i_{REF}(t)$$

The quality and dynamic properties of the compensation process depend mostly on the method used for calculating the reference-current parameters. Akagi *et al*'s theory of instantaneous reactive power [1] is commonly used to control power active filters. The authors believe that this theory does not fulfill the requirements of optimization of work in an energy source/receiver system. The general aim of optimization is to minimize the out-of-phase component of source current, reduce distortion of sinusoidal waveforms, and minimize active power losses in transmission of energy from source to receiver. To determine a current which would have such properties, we applied the variational methods of [10]. As a result, we obtained an expression describing an optimal source current (the target reference current) in the following analytical form:

$$i_{REF}(t) = {}_a i(t) = {}_e k(t) \cdot {}_e G(t) \cdot e(t) = A_{REF}(t) \cdot e(t)$$

Where: $e(t)$ is the voltage source, ${}_e G(t)$ the equivalent conductance in form: ${}_e G(t) = {}_a P(t)/E^2(t)$, where: ${}_a P(t)$ and $E(t)$ are instantaneous values of active power and rms voltage source [10]. The frequency and phase of the reference signal correspond to suitable values of the first harmonic of the voltage source, $e(t)$.

To effectively realize the whole control process, CM was divided into two sub-modules:

- the Identification Module (IM), which calculates frequency, ω_{REF} , phase, φ_{REF} , and amplitude, A_{REF} , of reference current, $i_{REF}(t)$,
- the Decision Module (DM), which performs these tasks:
 - shaping magnitude and phase characteristics of the active filter to obtain wideband transfer and high open-loop gain in the feedback loop. This is necessary to ensure a high degree of compensation of non-linear current and to work stably under conditions of a wide variety of load parameters,

- eliminating parasitic products of pulse-width modulation (PWM), used in generating i_{REF} from the feedback signal.

THE HARDWARE AND SOFTWARE

The prototype model of the compensator uses the Analog Devices ADDS-2106x-EZ-KIT microcomputer system, with the ADSP-21061 SHARC® floating-point digital signal processor. This high-performance system was needed because of the high calculating power required both by algorithms implemented within the Identification Module (IM) and in suitably shaping frequency transfer characteristics of the active filter. It is essential to assure the stability margin of all systems working in feedback closed loop under conditions of a wide range of load-parameter changes.

The evaluation system was developed with the addition of a universal analog and digital input/output module type ALS100, which had been designed by P.E.P. ALFINE as an extension of the ADDS-2106x-EZ-KIT. This module (Figure 2), designed for power-electronics applications, includes A/D and D/A converters, as well as PWM generators and a System Console (LCD & KBD). Communication with the host PC is established via an RS-232 port under control of a DSPHOST program.

Figure 2 shows the hardware and software structure of the Control Module. The main module of the control program was written in C language (ADDS-21000-SW-PC ver. 3.3), and time-critical procedures are written in Assembler.

The Control Module consist of:

- measuring resistors (R), in collaboration with transducers,
- an AD7864 four-channel, simultaneous-sampling A/D converter.
- A PWM generator that uses the ADMC201 motion coprocessor,
- System Console (SC),
- the Software Identification Module of reference current parameters (SIM),
- the Software Decision Module (SDM) collaborating with the Adder (Σ), which calculates current value of the error signal; i.e., the difference of reference and compensator current.

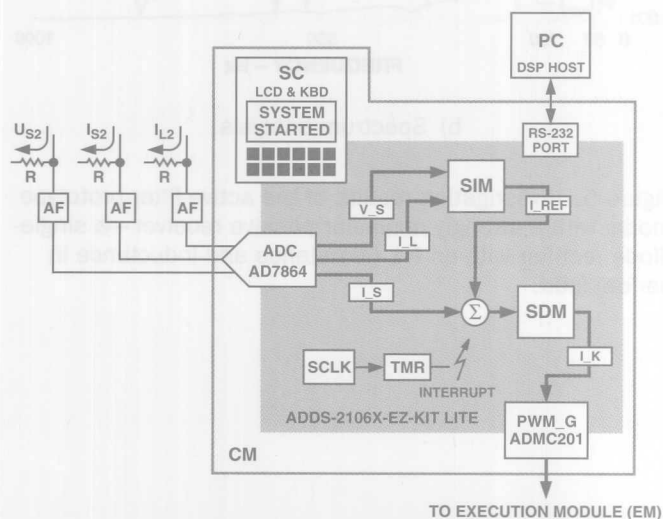


Figure 2. The hardware and software structure of Control Module (CM).

The SIM (Figure 3), consists of three independent blocks: *software-frequency-identifier* of the reference (SFI), *software-amplitude-identifier* of the reference (SAI) and *software synchronizer* of suitable values of the reference (SSYNC).

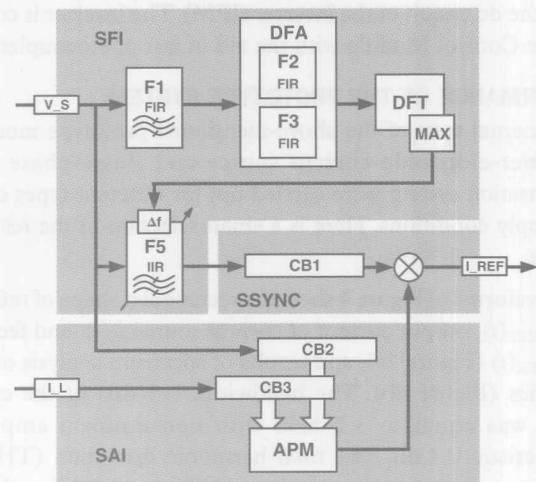


Figure 3. Software structure of Software Identification Module (SIM).

The SFI uses a mains-supply-voltage pre-filtration method, with the aid of a bandpass FIR filter (F1), to eliminate higher harmonics and increase noise immunity of the identification algorithm [14]. Next the signal is subjected to a Hilbert transformation to obtain its analytical form (complex signal in the time domain). It permits elimination of frequency products on the negative part of the frequency axis and decreases the identification time to under 12 ms. This is a short time in relation to the 20-ms (50-Hz) mains-voltage period of the present design, and would also be considerably shorter than the 16.7-ms period of 60-Hz systems [14]. The complex signal is subjected to a digital Fourier transform (DFT) to calculate its basic frequency. This is realized by the DFT and MAX blocks. Calculated in this way, the value of basic frequency serves next to control the tuned filter (F5), a high-Q, IIR-type filter. The F5 filter is in fact the reference current generator; its output signal frequency is equal to the mains-voltage frequency $u_s(t)$.

The amplitude of reference current is calculated within the SAI block, which is based on both load-voltage and load-current samples, stored within circular buffers CB2 and CB3.

A synchronization block, SSYNC, eliminates effects of different delay times, involved in calculations within the SFI and SAI blocks. Finally, the SSYNC connects suitable values of frequency and magnitude of reference current. The total time of identification and synchronization of the reference current generator (in this design) is about 18 ms.

The Decision Module is realized in the form of a 2nd order FIR filter with constant coefficients; its frequency transmittance model is given by following equation:

$$|T_{FK}(\Omega)| = \frac{1 + \cos(\Omega)}{2}$$

The basic condition of proper operation of the filter is that the system sample frequency is twice the PWM carrier frequency (in this system: 30 and 15 kHz).


The Execution Module is a power-electronics controlled current source, which uses a highly integrated Intelligent Power Module (IPM) type PM50RSA120 (Mitsubishi) and inductance coil, L_{INV} . This coil also limits parasitic products of the PWM.

The general source of energy for the current source is a capacitor within the dc circuit of the inverter (IPM). The inverter is coupled with the Control Module with the aid of fast photocouplers.


PERFORMANCE OF THE PROTOTYPE SYSTEM

Experimental tests of the above-mentioned prototype models of the power-electronic current source and single-phase active compensation system were carried out for different types of load and supply conditions. Here is a small selection of the results of the tests.

The waveforms of Figure 4 show the rectangular shape of reference signal $i_{REF}(t)$, output current of current source $i_S(t)$ and feedback signal $i_{FK}(t)$ (Figure 4a), and results of spectrum analysis of these quantities (Figure 4b). The bandwidth (-3 dB) of the current source was equal to 3.2 kHz with non-uniform amplitude characteristic 0.4 dB. The total harmonic distortion (THD) of output current within this band was 0.7%—and 0.2% within the 0.5-kHz bandwidth.




a) Waveforms of selected quantities.




b) Spectrum analysis.

Figure 4. Investigation results of current source prototype system for the case of rectangular shape of reference signal.

Figures 5 and 6 illustrate the workings of the complete active filter. The source of the distorted current (Figure 5) is a simple single-diode rectifier with an R-L-type load (resistor and inductor in series). It is a particularly unfavorable case, because it simultaneously generates a strongly distorted current with a dc component and reactive power. The waveforms of source voltage, u_S , and currents of load, i_L , power network, i_S , active filter, i_{INV} and reference signal i_{REF} , are shown in Figure 5a—and also results of spectrum analysis of selected quantities (Figure 5b). Figure 6 shows similar quantities for an RC-loaded 4-diode-bridge, the typical configuration of most consumer-electronics power packs.




a) Waveforms of selected voltage and current quantities.




b) Spectrum analysis.

Figure 5. Investigation results of the active-filter prototype model with a strongly nonlinear passive receiver—a single-diode rectifier with an R-L (resistance and inductance in series) load.



a) Waveforms of selected voltage and current quantities.



b) Spectrum analysis.

Figure 6. Investigation results of the active-filter prototype model with an RC-loaded 4-diode bridge.

As in the case of a current source, the active compensation's system of differential current provides good mapping of the reference signal, $i_{REF}(t)$, which is calculated in the Identification Module. The power network current is in the same phase as the waveform of power network voltage (because of compensation of so-called reactive power), and its higher harmonics values are considerably reduced. The THD value of active filter input current, $i_s(t)$, was under 1%.

CONCLUSION

We have shown here a system capable of achieving optimal compensation in real time by elimination of differential current, employing an applied power-electronic controlled current source using PWM. Included are a functional block diagram and a description of the working principle of the system, which is controlled by a digital signal processor. The measured results of tests carried out on the system for a variety of loads showed that the compensator was highly effective. It greatly reduced both the nonlinear distortion of the input current (THD < 1%) and the requirement for reactive power from the power source. The delay in identifying reference-signal parameters was about 12 ms (substantially less than one period of power-source frequency),

with total frequency identification error of 0.1%. In general all investigations results of the prototype model show very good mapping of the reference signal by the compensator system and considerable reduction of higher harmonics of source current.

REFERENCES

- [1] Akagi H., Kanazawa V., Nabae A.: *Generalized theory of the instantaneous reactive power in three-phase circuits*. Proc. of JIEE, IPEC Tokyo, 1983, 1375-1386.
- [2] Bayod Rujula A.A., Sanz Badia M.: *A new approach to harmonic compensation with hybrid active filters*. Proceedings of the 6th European Conference on Power Electronics and Applications, EPE'95, 1995, 1, 925-928.
- [3] Blajszczak G.: *Non-active power compensation using time-window method*. ETEP'92, Sept./Oct. 1992, 2, No. 5, 285-290.
- [4] Czarnecki L.S.: *Interpretacja, identyfikacja i modyfikacja własności energetycznych obwodów jednofazowych z przebiegami odkształconymi*. Monografia, ZN Pol. Sl., Elektryka, Z. 91, Gliwice 1984.
- [5] Fryze S.: *Moc rzeczywista, urojona i pozorna w obwodach o przebiegach odkształconych prądu i napięcia*. Przegląd Elektrotechniczny, 1931, No. 7, ss. 193-203.
- [6] Gwozdz M., Porada R.: *Eliminacja prądu różnicowego za pomocą kompensatora parametrycznego*. Proc. of XIX SPETO'96, Ustron, maj 1996, 383-386.
- [7] Gwozdz M., Porada R.: *Energoelektroniczny kompensator prądu różnicowego*. Materiały XX Seminarium z Podstaw Elektrotechniki i Teorii Obwodów, Gliwice-Ustron, 1997, 2, 341-344.
- [8] Koozehkanani Z.D., Mehta P., Darwish M.K.: *Active symmetrical lattice filter for harmonic current reduction*. Proceedings of the 6th European Conference on Power Electronics and Applications, EPE'95, 1995, 1, 869-873.
- [9] Pirog S.: *Three-phase active filter with sliding-mode control*. Proceedings of 7th International Power Electronic & Motion Control Conference and Exhibition, PEMC'96, Budapest, Hungary, 2-4 September, 1996, 1, 363-367.
- [10] Porada R.: *Minimal active current in circuits with real sources*. Proc. of 7th PEMC'96, Budapest, September 1996, 1, 405-409.
- [11] Pasko M.: *Dobór kompensatorów optymalizujących warunki pracy źródeł napięć jednofazowych i wielofazowych z przebiegami okresowymi odkształconymi*. Monografia, ZN. Pol. Sl., Elektryka, Z. 135, Gliwice 1994.
- [12] Siwczyński M.: *Metody optymalizacyjne w teorii mocy obwodów elektrycznych*. Inżynieria Elektrycz-na, Nr 183, Krakow 1995
- [13] Walczak J.: *Optymalizacja energetyczno-jakościowych właściwości obwodów elektrycznych w przestrzeni Hilberta*. Monografia, ZN. Pol. Sl., Elektryka, Z. 125, Gliwice 1992.
- [14] Gwozdz M., Porada R.: *Identification of basic frequency of periodical signals*. Proceedings of 7th International Power Electronic & Motion Control Conference and Exhibition, PEMC'96, Budapest, Hungary, 2-4 September, 1996, 1, 305-309. ▢

Ask the Applications Engineer—27

By Bill Englemann

SIGNAL CORRUPTION IN INDUSTRIAL MEASUREMENT

Q. *What problems am I most likely to run into when instrumenting an industrial system?*

A. The five kinds of problems most frequently reported by customers of our *I/O Subsystems (IOS) Division* are:

1. GROUND LOOPS

Ground loops are the bane of instrumentation engineers and technicians. They cause many lost hours troubleshooting obscure and hard-to-diagnose measurement problems. Do these symptoms sound familiar?

- Readings slowly drift even though you know the sensor is not changing.
- Readings shift when another piece of equipment is turned on.
- Measurements differ when a calibration device is connected at the end of an instrument cable instead of directly at the input.
- A 60-Hz sine wave is superimposed upon your dc measurement input.
- There are unexplained measurement equipment failures.

Any of these problems can be caused by ground loops—inadvertent flows of current through “ground,” “common” and “reference” paths connected to points at nominally the same potential. And all of these problems can be eliminated by *isolation*, the key signal-conditioning attribute we offer in all our signal conditioning series.

Sometimes separate grounding of two pieces of equipment introduces a potential difference and causes current to flow through signal lines. Why would this happen if they were both grounded? Because the earth and metal structures are actually relatively poor conductors of electricity when compared with the copper wires that carry power and signals. This inherent resistance to current flow varies with the weather and time of year and causes current to flow through any wires that are connecting the two devices. Many factory and plant buildings experience potentials of several tens or hundreds of volts. Appropriate signal conditioning eliminates the possibility of ground loops by electrically *isolating* the equipment. Signal conditioning will also protect equipment, rejecting potentially damaging voltage levels before entering the sensitive measurement system.

Isolation provides a completely floating input and output port, where there is no electrical path from field input to output and to power. Hence, there is no path for current to flow, and no possibility of ground loops.

Q. *How is this possible? How can we provide a path for the signal from input to output, without any path for current to flow?*

A. It's done by magnetic isolation. A representation of the signal is passed through a transformer, which creates a magnetic—not a galvanic—connection. We have perfected the use of transformers for accurate, reliable low-level signal isolation. This approach employs a modulator and demodulator to transmit the signal across the transformer barrier, and can achieve isolation levels of 2500 volts ac.

One of the most frequently encountered application problems involves measuring a low-level sensor such as a thermocouple in the presence of as much as hundreds of volts of ground potential. This potential is known as *common-mode voltage*. The ability of a high-quality signal conditioner to reject errors caused by common-mode voltage, while still accurately amplifying low-level signals is known as *common mode rejection (CMR)*. Our 5B, 6B and 7B Series signal conditioning subsystems provide sufficient common-mode rejection to reduce the impact of these errors by a factor of 100 million to 1!

2. MISWIRING AND OVERVOLTAGE

You know what happens when a cable from a sensitive data acquisition board is routed into another cabinet, or another part of the building—the input and output wiring terminals are grouped among hundreds of other terminals carrying diverse signals and levels: dc signals, ac signals, milli-voltage, thermocouples, dc power, ac power, proximity switches, relay circuits, etc. It's not difficult to imagine even a well trained technician or electrician connecting a wire to the wrong terminal. Wiring diagrams are often updated in real time with a red pen, as system needs change. Equipment gets replaced with “equivalents.” Sometimes power supplies fail and excess voltages are applied inadvertently. What can you do to protect your measurement system?

The answer lies in using rugged signal conditioning on every analog signal lead. This inexpensive insurance policy provides protection against miswiring and overvoltage on each input and output signal line. For example, the use of a 5B Series signal conditioner will provide 240-V ac of protection, even on input lines used to measure sensitive thermocouple signals, with levels in the millivolt range. You can literally connect a 240-V ac line across the same input lines used to measure the thermocouple, without any damage. The use of signal conditioning to interface with field I/O will protect all measurement and data acquisition equipment on the system side.

3. LOSS OF RESOLUTION

Resolution is the smallest change in the measurement that the analog-digital converter (ADC) system can detect and respond to. For example, if a temperature reading steps from 100.00° to 100.29° to 100.58°, as the actual temperature gradually increases through this range, the resolution (least-significant-bit value) is 0.29°. This would occur if you had a signal conditioner measuring a thermocouple with a range of 0° to +1200° and a 12-bit ADC. There are two ways to improve this (make the resolution smaller) and detect smaller changes - use a higher resolution ADC or use a smaller measurement range.

For example, a 15-bit plus sign ADC of the type used in our 6B Series would offer resolution of 0.037° on the 0 to 1200° range example, 8 times smaller! On the other hand, if you knew that most of the time the temperature would be in the vicinity of 100°, you could order from Analog Devices a thermocouple signal conditioner with a custom range, calibrated for the exact thermocouple type and temperature measurement range. For example, a custom-ranged signal conditioner with a span of +50° to +150° would offer resolution of 0.024° with a 12-bit ADC, a big improvement over the 0 to 1200° range.

4. MULTIPLE SIGNALS DON'T ALL HAVE THE SAME PROPERTIES

This can pose quite a challenge to traditional industrial measurement approaches where 4, 8 or even 16 channels are dedicated to interfacing to the same signal type. For example, let's say you need to measure two J thermocouples, one 0 to +10 V signal, four 4-20 mA signals and two platinum RTDs (resistance temperature detector). You can either buy individual transmitters for each channel and then wire them all into a common 4-20 mA input board, or use a signal conditioning solution from Analog Devices that is configured channel-by-channel, but is also integrated into a simple backplane subsystem.

These subsystems incorporate all connections for input, output and field wiring, as well as simple connections for a dc power supply. They offer a choice of output options: 0 to +5 V, 0 to +10 V, 4-20 mA and RS-232/485, and more! Input and output modules are mix-and-match compatible on a per-channel basis and hot-swappable for the ultimate flexibility.

5. ELECTRICAL INTERFERENCE

Today's industrial factories and plants contain all kinds of interference sources: engines and motors, fluorescent lights, two-way radios, generators, etc. Each of these can radiate electro-magnetic noise that can be picked up by wiring, circuit boards and measurement modules. Even with the best shielding and grounding practices, this interference can show up as noise on the signal measurement. How can this be eliminated? By providing high noise rejection in the signal conditioning subsystem.

Lower-frequency noise can be eliminated by choosing signal-conditioning subsystems with excellent common mode and normal mode rejection. Common mode noise present on both the plus and minus inputs can be seen when measuring either the plus or minus input with respect to a common point like ground. Normal-mode noise is measured in the difference between the plus and minus inputs. A typical common mode rejection specification on our signal conditioning subsystems is 160 dB. This log scale measurement means that the effect of

any common mode voltage noise is reduced relative to signal by a factor of 10^8 , or 100 million to 1!

Very high frequency noise in the radio frequency bands can cause dc offsets due to rectification. It requires other approaches, including careful circuit layout and the use of RFI filters such as ferrite beads. The performance measures are indicated by our compliance with the EN certifications for electromagnetic susceptibility popularized by the CE mark requirements of the European community. A typical application where this is important would be where a two-way radio is used within a few feet of the input wiring and signal conditioning subsystem. It is necessary to reject measurement errors whenever the radios are transmitting. Good panel layout practice and the use of signal conditioning will ensure the best accuracy in these noisy environments.

CONCLUSION

Q. What are some good installation and wiring practices?

A. Here are a few suggestions. You may also want to take a look at "Design Tools" and the Analog Devices book, *Practical Analog Design Techniques*, available for sale in hard copy and free on the Web.

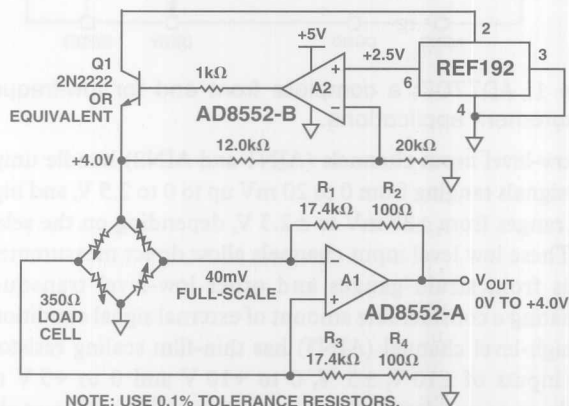
- Avoid installing sensitive measuring equipment, or wire carrying low level signals, near sources of electrical and magnetic noise, such as breakers, transformers, motors, SCR drives, welders, fluorescent lamp controllers, or relays.
- Use twisted pair wiring to reduce magnetic noise pickup. Look for 10 to 12 twists per foot.
- Use shielded cable with the shield connected to circuit common at the input end only.
- Never run signal-carrying wires in the same conduit that carries power lines, relay contact leads or other high-level voltages or currents.
- In extremely high interference environments, mount signal conditioning and measurement equipment inside grounded and closed metal cabinets. ▢

New-Product Brief

AMPLIFIERS

Zero-Drift, Single-Supply, Rail-to-Rail Input/Output Operational Amplifiers

The AD8551, AD8552, and AD8554 are single, dual, and quad high-accuracy op amps, featuring 1 μ V offset voltage, 5 nV/ $^{\circ}$ C drift, 130 dB PSRR, 140 dB CMRR, 145 dB Voltage Gain, and 10 pA bias current. They operate on a single 2.7- to 5-V supply and draw only 750 μ A per amplifier. With 20-bit accuracy, the AD855X is perfectly suited for the most demanding applications in medical, industrial and automotive instrumentation, battery management and a wide range of temperature, current, voltage, pressure and other sensor applications. They are packaged in 8 or 14 lead MSOP, TSSOP, and SOIC packages, and operate over the -40° C to $+125^{\circ}$ C temperature range. Price (1000s) for the singles/duals/quads is \$1.14/\$2.11/\$4.05.



A +5 V Precision Strain-Gage Amplifier

Process Signals from Millivolts to ± 10 V Directly with a Versatile Single-Supply 3/5-V Charge-Balancing A/D Converter

by: Albert O'Grady,
Applications Engineer, General-Purpose Converters

The AD7707, shown in Figure 1, is an oversampling A/D converter with on-chip digital filtering. It is intended for the measurement of wide-dynamic-range, low-frequency signals, such as those encountered in industrial control or process-control applications. It contains a multiplexer, programmable-gain amplifier, sigma-delta (charge-balancing) ADC, a calibration microcontroller with on-chip static RAM, a clock oscillator, a digital filter and a bi-directional serial communications port. The AD7707's three analog input channels include two low-level quasi-differential inputs and one high-level single-ended channel.

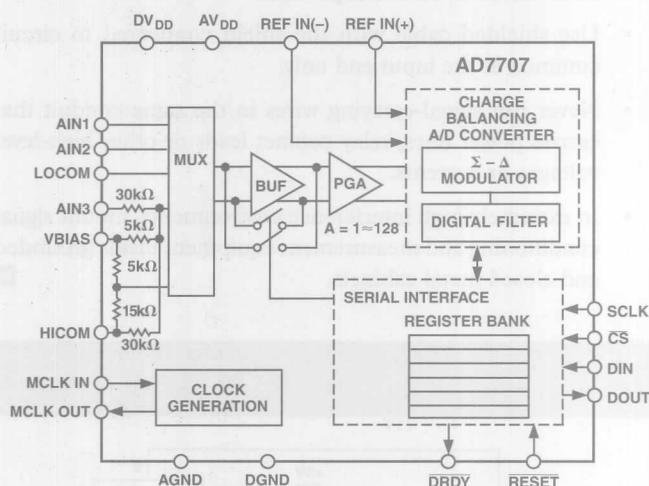


Figure 1. AD7707: a complete front end for low-frequency measurement applications.

The low-level input channels (AIN1 and AIN2) handle unipolar input signals ranging from 0 to 20 mV up to 0 to 2.5 V, and bipolar input ranges from ± 20 mV to ± 2.5 V, depending on the selected gain. These low level input channels allow direct measurement of signals from strain-gauges and other low-level transducers, eliminating a considerable amount of external signal conditioning. The high-level channel (AIN3) has thin-film scaling resistors to allow inputs of ± 10 V, ± 5 V, 0 to +10 V and 0 to +5 V to be directly accommodated without requiring split supplies or charge pumps.

The AD7707's sigma-delta conversion technique realizes up to 16 bits with no missing codes. The selected input signal is applied

to a proprietary programmable-gain front end with an analog modulator. The modulator output is processed by an on-chip digital filter. The first notch of this digital filter can be programmed via an on-chip control register to allow adjustment of the filter cutoff and output update rate.

Featuring a serial interface that can be configured for three-wire operation, the AD7707 is ideal for use in smart microcontroller- or DSP-based systems. Gain settings, signal polarity and update-rate selection can be configured in software using the input serial port. The device contains both self-calibration and system-calibration options; these allow gain and offset errors, either on its own part or in the system as a whole, to be corrected.

The AD7707 operates from a single 3-V (2.7 to 3.3) or 5-V (4.75 to 5.25) supply. Consuming less than 1 mW with 3 V supplies, using a 1-MHz master clock, it is ideal for use in low-power systems; in *standby* mode it draws less than 8 μ A. The AD7707 is available in both a 20-lead 0.3"-wide small-outline (SOIC) package and a low profile 20-lead TSSOP. The combination of low power and small footprint makes the device ideally suited for use in field equipment. The versatility of its high-resolution-and-accuracy high- and low-level input voltage channels comes at a low cost (a major factor in selecting data acquisition components), just \$4.46* (1000s).

The AD7707 is well-suited to applications in data acquisition requiring true bipolar input capability in a single supply system. These opportunities include low-level direct transducer interface applications, such as those found in pressure- and temperature measurement applications, smart-valve/actuator control systems, smart transmitters and chart recorders.

SMART-VALVE/ACTUATOR CONTROL

In this example (Figure 2), the desired valve set point is established remotely in the control room and communicated to the AD7707 with low sensitivity to noise, using a 4-to-20-mA current-loop. This control signal is converted by the AD7707 to digital, along with the valve-position signal, and they are compared by a microcontroller, which operates the valve actuator, closing the loop. Both devices operate at low voltage from single supplies.

The 4-to-20-mA valve-position control signal is translated from digital in the control room by an AD420 DAC, with 4-to-20-mA output. The current, delivered to the valve's vicinity via a noise-rejecting twisted pair, is sensed by the AD7707's low-level input channel, and the resulting voltage is converted to digital. The valve position is monitored using a high-quality servo-potentiometer with ± 10 -V output range and applied as a high-level analog input to the AD7707. The controller, the valve, and the AD7707 thus form a closed loop control system.

Using the high level input channel on the ± 10 -volt input range, peak to peak resolutions of 16 bits can be achieved with update rates of up to 60 samples per second. The low-level input channels also provide 16-bit peak to peak resolution when operating at a gain of one, using a 125-ohm sense resistor on the 4-to-20-mA loop. With a 10-Hz update rate, filter notches are placed at both 50 and 60 Hz, simultaneously rejecting components at both frequencies—a key requirement in industrial applications.

This, and many other applications, benefit from the use of an ADC operating from a single supply, ability to deal with both high and

*Prices are recommended resale prices (U.S. Dollars) FOB U.S.A. Prices are subject to change without notice. For specific price quotations, get in touch with our sales offices or distributors.

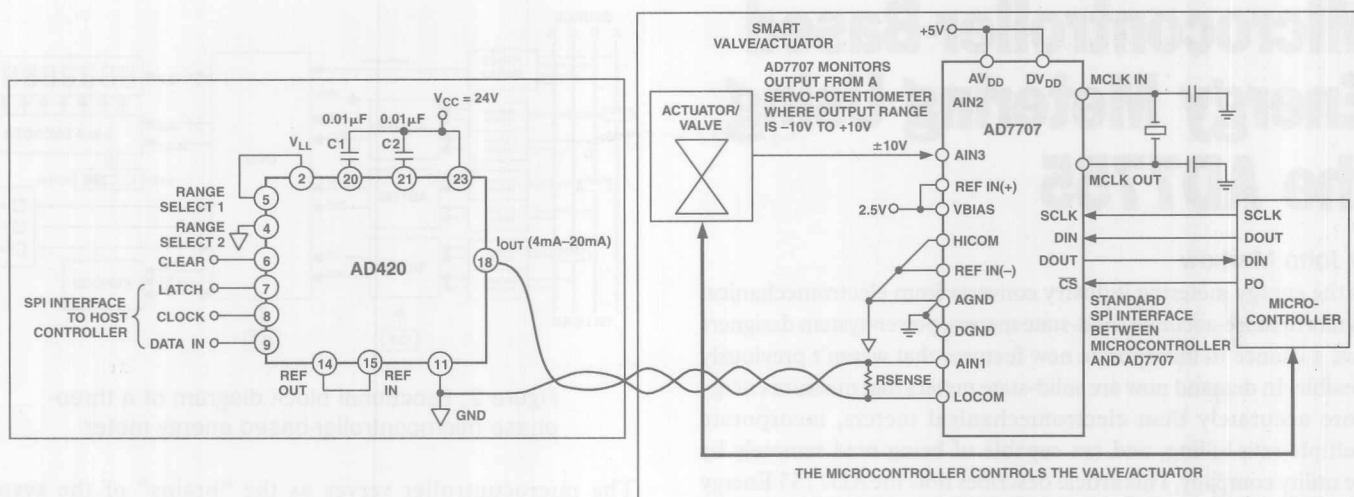


Figure 2. Smart-valve/actuator control using the AD7707.

low level analog input voltages, and high-resolution analog-to-digital conversion at low cost. The use of sigma-delta architecture for the analog-to-digital function provides a degree of immunity to noisy environments. This, combined with a programmable gain amplifier, a digital filter and calibration options, makes the ADC ideal for use in industrial and process control applications. The AD7707, with its excellent system noise performance and lack of need for high-quality external capacitors, manifestly provides far more system level functionality than off-the-shelf integrating ADCs.

CHART RECORDERS

Another area where both high- and low level input channels are usually required is in chart recorder applications. Circular chart recorders generally have two requirements. The first would utilize the low-level input channels of the AD7707 to measure directly inputs from thermocouples, RTDs, and pressure sensors. The second requirement, to be able to measure dc input voltage ranges

up to ± 10 V, would be provided by the high-level input channel without needing external signal conditioning. These requirements can be satisfied by the AD7707's combination of input channels without external signal-conditioning components, split supplies, or charge pumps. It also meets the key requirement of low-power operation in portable data acquisition equipment, which is also a feature of the AD7707.

Brief Performance Summary

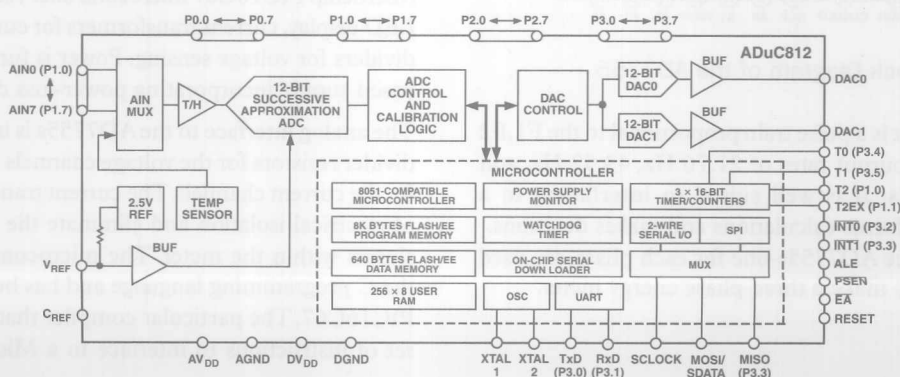
Input	AIN3	AIN1, AIN2	AIN1, AIN2
Range	± 10 V	± 20 mV	± 20 mV
Update Rate	60 Hz	10 Hz	60 Hz
Resolution	16 Bits	14.5 Bits	13 Bits
Remarks	G = 2	Unbuffered; 5-V Supplies; 2.5-V Reference	

New-Product Brief

Microconverter™ Complete 12-Bit Data-Acquisition System with Microcontroller and Flash/EEPROM

The ADuC812 combines conversion, processing, and memory (8-channel 12-bit A/D converter and two 12-bit voltage-output DACs; an 8051 microcontroller; and Flash data and program memory) in a single monolithic chip to form a general-purpose "plug-and-play" data-acquisition system. Truly flexible, it can acquire analog and digital data and handle calibration, linearization, signal-level translation, communications, and datalogging.

It can be used as a user-programmed front end; and its Flash/EE memory can be rewritten without requiring a high-voltage supply. Among its features are an on-chip reference, SPI, 12C, and UART serial interface, watchdog timer, and power-supply monitor/reference sensor. Typical applications include intelligent sensors (IEEE 1451.2-compatible), battery-powered instrumentation, transient-capture systems, and data-acquisition systems. It will operate on 3- or 5-V supplies, at temperatures from -40 to $+85^{\circ}\text{C}$, is housed in a 52-lead PQFP. Price in 100s is \$8.98.



Microcontroller-Based Energy Metering Using the AD7755

by John Markow

As the energy metering industry converts from electromechanical meters to more-accurate solid-state meters, power-system designers have a chance to incorporate new features that weren't previously possible. In demand now are solid-state meters that measure energy more accurately than electromechanical meters, incorporate multiple-rate billing, and are capable of being read remotely by the utility company. This article describes how the AD7755 Energy Meter¹ integrated circuit could be used in three-phase energy metering with power outage detection and measurement backup, and remote, automated, multiple-rate metering.

The AD7755 is an accurate (to 0.1%) single-phase energy-measurement IC. It accepts a pair of voltage inputs that represent the voltage and current of a power line. Internally, these signals are converted to the digital domain with oversampling A/D converters. A fixed-function digital signal processor continuously multiplies the two signals; their product is proportional to instantaneous power. After being low-pass filtered, the digital signal is then converted to a frequency—scaled according to selectable settings—to generate frequency outputs at terminals F1, F2, and CF. The signals at F1 and F2 can be used to drive an electromechanical counter (typically at full-scale rates from 0.5 to 5 Hz), while the higher-frequency CF signal is suitable for calibration. The *frequency* (or rate) of the pulse outputs is proportional to the instantaneous *real power* being monitored by the meter. Accordingly, in a given interval, the total *number* of pulses generated at these outputs is proportional to the *energy* transferred to the load. A reverse-polarity logic signal indicates when the measured instantaneous power goes negative (i.e., the load is returning net power to the line).

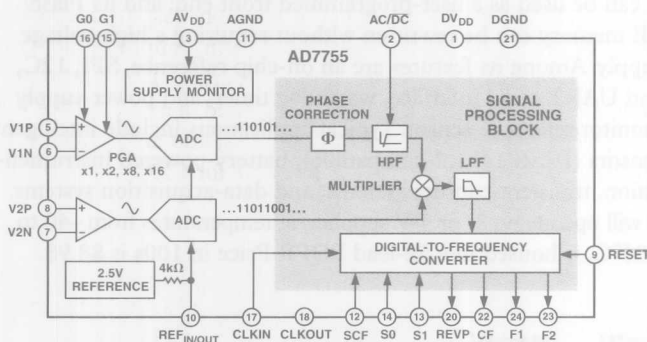


Figure 1. Block Diagram of the AD7755

The CF frequency output is a pulse train proportional to the F1, F2 outputs, with full-scale output rates of 21.76 Hz, 43.52 Hz, and 5.57 kHz, for ac inputs. It is well suited to interfacing to a microcontroller that performs calculations and makes decisions. Figure 2 shows how three AD7755s—one for each phase—are used with a microcontroller to make a three-phase energy meter.

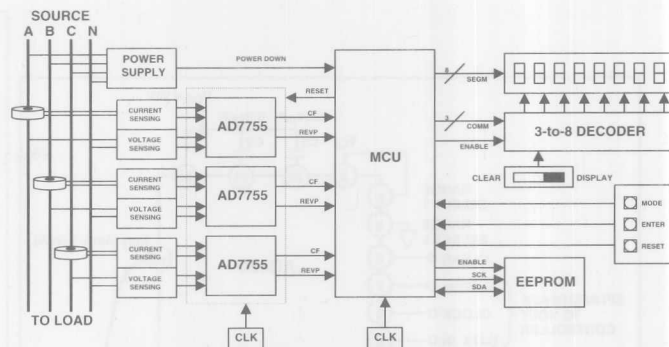


Figure 2. Functional block diagram of a three-phase microcontroller-based energy meter

The microcontroller serves as the "brains" of the system, performing all the required housekeeping tasks and interacting with the other components—the energy meter ICs, the power supply, the EEPROM, the display, and buttons to operate the meter—to view energy or power, calibrate the phases, or clear the reading. Besides low cost, the basic microcontroller requirements are:

- **Sufficient I/O to drive the display.* If an LCD display is used a driver is required. If one is not incorporated into the MCU, an LED display can easily be controlled with a 3-to-8 decoder.
- **Interrupts.* To avoid missing any energy-indicating pulses, the system can be configured to trigger interrupts in the MCU. A power supply monitor can generate an MCU interrupt when it has detected a brownout condition and initiate an emergency energy measurement backup.
- **EEPROM Serial Interface.* A simple serial interface can be created using only two or three I/O lines. An MCU with a built-in serial interface makes the design even easier.
- **Timers.* There are two main time intervals that need to be maintained. First, a display update rate must be set at about 2 seconds. Also, if an LED display is used, a timer must cycle through the digits at a sufficient rate to minimize on flicker. Additionally, the calibration routine must be carefully timed, but can be implemented with interrupt postscales.

As an added feature, a second serial interface could be used to communicate with a host system for remote/automated metering. Also, either an external or internal clock could be used to implement multi-rate metering.

Reference Design: A three-phase energy-meter reference design (Figure 2) has been implemented to demonstrate how multiple AD7755s can be interfaced to a microcontroller. It uses a Microchip PIC16C67 microcontroller², serial EEPROM, an 8-digit LED display, current transformers for current sensing, and resistor dividers for voltage sensing. Power is furnished by a transformer-based supply incorporating power-loss detection.

The analog interface to the AD7755s is instrumented with voltage divider resistors for the voltage channels and current transformers for the current channels. The current transformers provide a degree of electrical isolation and eliminate the need for current-sensing shunts within the meter. The microcontroller code is written in the C programming language and has been programmed into the PIC16C67. The particular compiler that was used also includes a set of instructions to interface to a Microchip serial EEPROM,

¹www.analog.com

²www.microchip.com

which stores energy measurement digits and the calculation limits obtained in calibration mode. The display consists of 8 LED digits multiplexed by a 3-to-8 decoder. The power supply uses three transformers, a rectifier, and a regulator to convert 220-V, 3-phase ac to a dc voltage capable of powering the meter even if one phase goes out. A reference design data sheet/application note is available³.

Since the meter determines the cost of energy to the user, the most important requirements of an energy meter are reliability and accuracy over time. The energy is measured in a fairly simple way—by pulse-limit comparison. In this method, the microcontroller counts the number of pulses on a phase until the total reaches a calibrated limit. At this point, the energy reading is incremented by the smallest unit within the range of the display (in this case, 0.01 kWh). This technique implies that the display register need only be updated when necessary and also avoids complex numerical operations that could make the meter operate inefficiently.

The maximum output frequency at CF in the slow mode is 43.52 Hz, or 156,672 pulses/h. Allowing for headroom, a 220-volt, 60-ampere system could be calibrated so that 0.01 kWh of measured energy produces approximately 100 pulses on CF. Calibration is done in the high-frequency mode, in which the maximum frequency on CF is 128 times faster, or 5.57 kHz. During calibration, a fixed power is set on the line at a value that yields 1/128 kWh over the calibration interval. Because of the scaling factor of 128, the number of pulses counted during the calibration time is equivalent to 1 kWh when the AD7755 is returned to the low-frequency mode.

Suppose 10287 pulses had been counted during this interval. Then the display would have to be incremented by 0.01 kWh every

102.87 pulses, rather than every 100 pulses. This fractional-N count could be accomplished in several ways. For example, during the time required to advance the display by 100 increments (i.e., to accumulate 1.00 kWh in 0.01 kWh steps), 13 of those steps could be produced by a count of 102 pulses, and the other 87 steps would require a count of 103 pulses. An alternative that has been in use is, starting with 102 pulses per step, to add 8 pulses on every 10th step, and 7 more pulses on every 100th step.

The high frequency setting on the AD7755s yields better results in a shorter calibration time (about 30 seconds). Due to possible variations in the system, and to historical practices, each phase of a three-phase meter is calibrated independently.

Conclusion: A basic solid-state meter design like that described above will most likely be more accurate, more reliable, and cheaper than electromechanical meters, and will allow for added features that benefit both the customer and utility company. In the near future, the utility company will monitor your energy consumption remotely and bill you according to peak- and off-cycle usage—or even (in the Motorola⁴ system) allow you to keep track of your usage. Not only will the measurement be more accurate, but, implemented over an entire electrical network, solid-state metering allows for more efficient energy management.

References

Three-Phase Energy Meter reference design data sheet (REF-AD7755-3). Analog Devices.

AD7755 Energy Meter IC with Pulse Output data sheet. Analog Devices.

Daigle, Paul. "All Electronic Power and Energy Meters," *Analog Dialogue*. Volume 33, Number 2, February, 1999. ▀

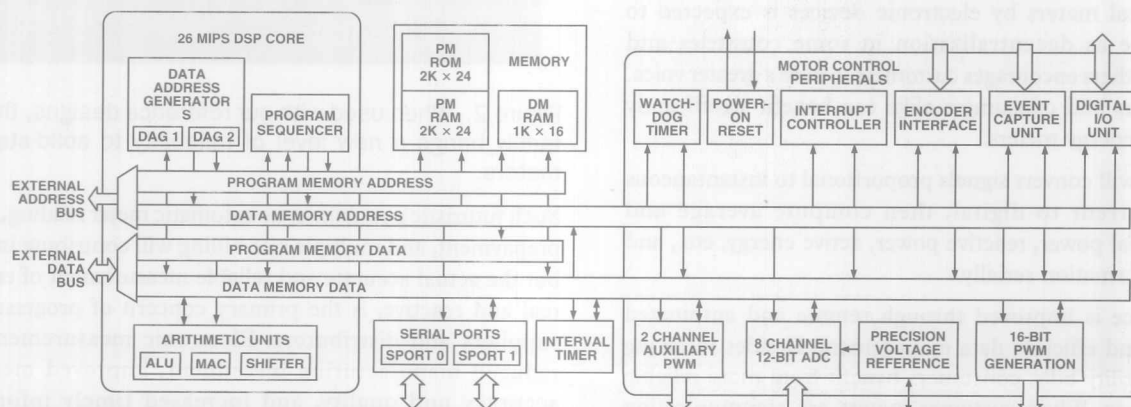
New-Product Brief

MOTOR CONTROLLER

Motor Controller, DSP-Based Single-Chip

The *ADMC401* is a programmable DSP-based motor controller. It uses a 16-bit fixed-point ADSP-2171 cell, capable of processing 26 MIPS (million instructions per second)—and multiple instructions per cycle—to deliver the highest processing power, precision and performance for complex motor control algorithms. The single-chip motor controller integrates the DSP, a high-resolution, 8-channel simultaneous-sampling 12-bit analog-to-digital converter, and peripherals on one integrated circuit. The

ADMC401's motor control peripherals include an encoder interface, pulse width modulation generator, power-on-reset, and programmable digital I/O, which help reduce device size and cost, without sacrificing performance. In addition, easy-to-use Analog Devices DSP tools, extensive application support, and DSP programmability speed user time-to-market. The *ADMC401* operates on *V*_{cc} from -40 to °C, and is housed in a 144-lead LQFP. It is priced below \$18 in OEM quantities, and an evaluation kit (*ADMC401-ADVEVALKIT*) is available for \$395.



³www.analog.com/industry/energymeter/index.html

⁴www.mot.com/GSS/SSTG/ISSPD/URM/index.html

All-Electronic Power and Energy Meters

By Paul Daigle, Product Manager

The AD775x family of energy-measuring integrated circuits accepts voltage inputs representing local voltage and current in an electrical power system and converts them to digital using oversampling A/D converters. An on-chip digital processor continuously computes the product of the two signals, which is proportional to instantaneous power. Input conditioning, filtering, further processing, and other features, which are specific to each type within the AD775x family, provide metering solutions for a variety of power-system applications.

For example, the most general-purpose member of the family, the AD7750 (Figure 1), low-pass filters the computed product, then uses digital-to-frequency conversion to output a complementary pair of pulse trains of frequency proportional to the instantaneous real power—for driving a counter or two-phase stepping motor—plus a higher-frequency output, suitable for calibration and test.

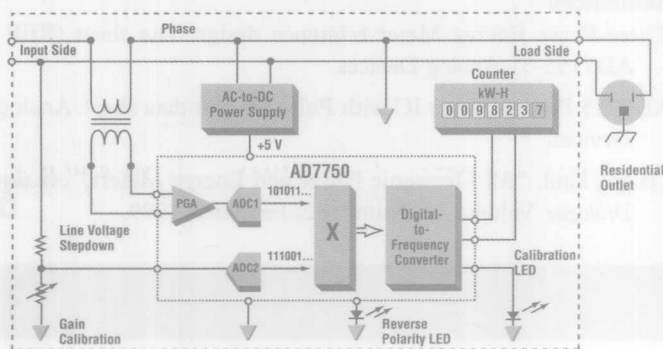


Figure 1. Functional block diagram of the AD7750 in a typical application measuring single-phase real power. Progress in power measurement.

Before continuing with a discussion of system applications and the features they give rise to, let us consider recent progress in power measurement in the electrical industry.

Electrical utilities, like many other heavy industries, have greatly increased their familiarity with sophisticated electronics in recent years. For example, the rate of replacement of long-used electromechanical meters by electronic devices is expected to quickly increase as decentralization in some countries and deregulation in others encourages customers to have a greater voice. Energy producers and consumers alike can benefit significantly from electronic energy meters.

A typical meter will convert signals proportional to instantaneous voltage and current to digital, then compute average and instantaneous real power, reactive power, active energy, etc., and transmit the information serially.

Customer service is improved through remote and automated meter reading and efficient data management. Besides receiving more credible utility bills, consumers benefit from more reliable power distribution. When customer meters are communicating through a network, power outages can be detected, identified, and corrected more quickly.

If the required ratio of peak power to average power in a system is reduced, the consequent reduction in required generating capacity will reduce environmental disturbance and pollution. The incentives provided by multiple-rate billing will help to greatly reduce peak usage despite population growth. Distribution cleanliness is maintained by monitoring the power-quality pollution (e.g., excessive reactive power, nonlinear loads, dc offsets) imposed by individual consumers. Consumers can benefit from lower electricity bills with the installation of smart card controlled energy meters that lower operational costs of providing service, reading meters, and processing data.

Electronic meters can compute power accurately irrespective of phase shifts and waveform distortion due to nonlinear loads; also, electromechanical meters are not able to accurately measure energy in the presence of phase-fixed load regulation schemes popular on distribution networks. Electronic measurements are thus more robust and accurate under these conditions.

Granted that electronic energy meters have outperformed the electromechanical meter in terms of functionality and performance, how do they stack up in cost and reliability? Two thumbs up! The entry to this field of companies like Analog Devices, with its excellent reputation for supplying analog, digital, and mixed-signal integrated circuits in large quantity for military, aerospace, and high-volume consumer products, promises the successful marriage between high reliability and low cost electronics that the industry has been waiting for. Recognizing the cost constraints of single-phase energy meters, ADI has identified an opportunity to help meter manufacturers meet their volume requirements, while reaching their cost targets and alleviating their reliability concerns.



Figure 2. When used with our reference designs, the AD775x family brings a new level of reliability to solid-state energy meters.

Such futuristic possibilities as automatic meter reading, smart card prepayment, and multiple rate billing will contribute importantly, but the actual accurate and reliable measurement of energy, both real and reactive, is the primary concern of progressive energy suppliers and distributors. Electronic measurement leads to reduced manufacturing investment, improved measurement accuracy and quality, and increased timely information, a combination of benefits that go well beyond the traditional rotor-plate energy meter design.

DSPs and Microcontrollers

The first attempts at electronic energy meters derived power by multiplying current and voltage in the analog domain, but the linearity over temperature and time proved to be no better than electromechanical meters. The stability, linearity, and accuracy provided by automatic error detection/correction of digital calculations has already swept across the communications industry and now has arrived at the door of electrical power metrology. Digital signal processing (DSP)-based products perform multiplication and other calculations on current- and voltage signals that have been digitized with on-chip analog-to-digital converters (ADCs). Processing the signals digitally provides stable and accurate calculations over time despite variations in the environment.

Although programmable DSPs are widely available at low cost and offer a degree of flexibility, what may turn out to be the most cost-effective form of processing electrical power measurements involves the use of a low-cost fixed-function (embedded) DSP, with on-chip A/D converters, for measurement and computation—and an associated microcontroller to handle programming tasks and simple calculations for communications and display. The DSP is continuously converting, sampling, and computing instantaneous and average power.

For example, during the past year many different energy meter designs have been manufactured using 4-MHz, 4-bit

microprocessors. Such *microcontrollers* allow a limited degree of configurability while managing some house keeping functions, such as data encryption and demodulation, time stamping for multiple rate billing, and energy-delivery intelligence (power outage detection, remote disconnect, prepayment, load management). The microprocessor allows users to select the level of service they want, and the utility can remotely configure individual meters.

Standard Products

The growing family of standard products designed for energy measurement not only eliminates the high manufacturing investment associated with electromechanical energy meters; it also greatly reduces the need to develop ASICs (application-specific integrated circuits). Standard products incorporate the solutions of problems common to multitudes of different customers at a lower shared cost. Factors that manufacturers of electronic energy meters should consider when seeking to optimize overall cost-effectiveness of measurement choices include accuracy, hardware, software, development costs, time to market, and ease of implementation.

The following table shows how Analog Devices's growing family of fixed-function DSPs addresses the wide variation of system considerations worldwide. The choice of component from the family depends on the type of meter that is required for a given system.

Integrated Power Meter ICs

	AD7750	AD7751	AD7755	AD7756	Single Phase VA & VAR	Three Phase VA & VAR
PHASES						
Single Phase, 2-Wire	•	•	•	•	•	
Two Phase or Single Phase, 3-Wire	•		•	•	•	
Three Phase, 3- or 4-Wire (Wye or Delta Load)			•	•	•	•
INTERFACE						
Micro-ohm Shunts & Current Transformers		•	•	•	•	•
Milli-ohm Shunts & Current Transformers	•	•	•	•	•	•
OUTPUT						
High Frequency Pulse	•	•	•	•	•	•
Real Power	•	•	•	•	•	•
Low Frequency Complementary Pulse	•	•	•			
External Calibration	•	•	•			
Internal Calibration				•	•	•
Fault Tolerant Billing		•				
Zero Crossing Logic Output (Frequency)				•	•	•
Interrupt Request Output				•	•	•
Serial Port Interface				•	•	•
Apparent & Reactive Power, Voltage, Current					•	•
PACKAGE OPTIONS						
20-PDIP and 20-SOIC	•		•	•		•
24-PDIP and 24-SSOP		•	•		•	
SAMPLES	NOW	NOW	NOW	Aug '99	2000	Nov '99

NOTE: The first few AD775x standard products—designed to directly drive a stepper motor counter—precede a series of products with a serial port interface for bidirectional communication with a microprocessor. Analog Devices, Inc. will continue to assist in driving the cost of energy meters down by addressing the costs of the power supply, current transducer, oscillator, and external gain calibration. Ultimately, a highly integrated product can be designed to meet the aggressive cost targets while maintaining a great deal of functionality by close cooperation with both meter manufacturers and utilities. Contact our Power Measurement Group via email to begin developing a relationship.

The first in the family, the *AD7750*, is designed to directly drive a stepper motor counter to integrate power to energy. In terms of cost, the stepper motor counter is popular in developing countries because it is a practical way to build an inexpensive solid state energy meter. When power is lost, the counter simply stops turning. Other solutions, like light-emitting diode (LED) displays or liquid-crystal displays (LCDs), require high pin count drivers and a method to store the reading during power loss.

The next series of products have a serial port interface for bi-directional communication with a microprocessor. Products available soon will assist in driving the cost of energy meters down by addressing the costs of the power supply, current transducer, oscillator, and external gain calibration. Ultimately, close cooperation with meter manufacturers and utilities will lead to a highly integrated device designed to provide a great deal of functionality while meeting aggressive cost targets.

The *AD7750* integrates two 16-bit analog-to-digital converters and the digital signal processing logic necessary to measure electrical energy. With the exception of the analog circuitry in the A/D converters and the reference circuit, all other signal processing (e.g., multiplication and filtering) is carried out in the digital domain. This approach provides superior stability and accuracy over extremes in environmental conditions and over time. The sigma-delta converters, operating at an oversampled rate of 900-kHz, which simplifies antialiasing, digitize the voltage signals from current and voltage transducers. The current channel has wide dynamic range and programmable gain to deal with direct connection to a variety of current-to-voltage transducers, which generally have low terminal voltage. A high-pass filter removes any dc from the current channel, eliminating inaccuracies that offset voltages might introduce into the calculation of real power.

Real power is calculated from the instantaneous power signal, which is generated by multiplying the current and voltage signals. A

high-pass filter can be switched into the signal path of the current channel to remove any offsets. Low-pass filtering reduces the line-frequency harmonics and extracts the real power (in other words, dc) component. This approach calculates real power correctly, even with non-sinusoidal current and voltage waveforms and any power factor. The digital signal processing, (multiplication, filtering, etc.) ensures high stability over temperature and time.

The chip also contains two digital-to-frequency converters; one has a low-frequency output, the other, a high-frequency output. In both cases, the output pulse rate of the digital-to-frequency converters varies with value of real power dissipated over time. The chip offers a range of output frequencies, selectable by the designer, to accommodate most meters. The low-frequency output, because of its long accumulation time between pulses, has a frequency that is proportional to the average real power. The high-frequency output, with its shorter accumulation time, is proportional to the instantaneous real power. As a result, the high-frequency output is useful for calibrating the meter under steady load conditions.

Energy Metering IC with On-Chip Fault Detection

The *AD7751* is an accurate fault-tolerant electrical energy measurement IC intended for use in two-wire distribution systems. The part incorporates a novel fault-detection scheme, which both warns of fault conditions and allows the *AD7751* to continue accurate billing despite a fault event. It does this by continuously monitoring both the phase and neutral (return) currents. A fault is indicated when these currents differ by more than 12.5%, and billing is continued using the larger of the two currents.

Energy Metering IC with Pulse Output

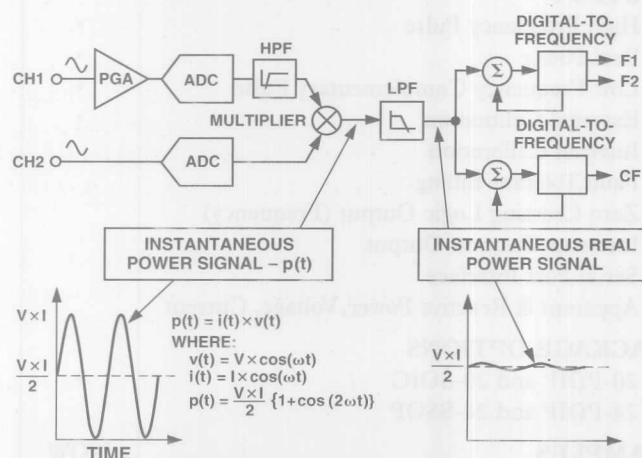
The *AD7755* core is pin-compatible with the *AD7751* but it does not include the fault-tolerant billing feature. It is also designed so that it can be used in systems with more than 2 wires, including 2- and 3-phase systems. ▶

New-Product Brief

Energy Meter

Single Phase, Fault Tolerant Energy Measurement IC

The *AD7751* is a data converter IC used to make the lowest cost, single phase energy meter that is fault tolerant. It has two current channels and one voltage channel, each with a dedicated sigma-delta A/D converter. Multiplication of the larger of the two currents and the voltage is done in the digital domain. The result is real power. Average power can be accumulated by a stepper motor counter to obtain and display energy in kilowatt hours. A high frequency output provides instantaneous real power for calibration and communication with a microprocessor. The *AD7751* typically has less than a 0.3% error in measurement of true energy over a 1000 to 1 dynamic range, surpassing the requirements of the IEC1036 standard. It is available in 24-lead DIP and SSOP packages, and operates over the -40°C to $+85^{\circ}\text{C}$ temperature range. Price (1000s) for the ARS/BRS versions is \$2.83/\$3.40.



MAKE vs. BUY?

When Should I Reinvent the Wheel?

by Rich Ghiorse and Mike McKeon

The thorny choice between making and buying common circuit functions is faced by many system designers—and their companies—on a regular basis. This choice is also of great interest to companies (such as Analog Devices) that design, manufacture, and market input/output (I/O) functions. These companies have successfully served the industrial systems market by offering useful and reliable functions with guaranteed performance at modest prices.

Industrial I/O functions are complete signal conditioning solutions for companies whose primary business is design and assembly of large control systems. Typical functions are used for measuring temperature and pressure, interfacing with strain gauge signals, and controlling actuators. I/O products serving this market have to be reliable, endure harsh environments, meet agency approvals, and be cost-effective.

Should companies have to reinvent common I/O functions (the wheel)? Or should they purchase the solution from a vendor like ADI? Those who choose the option of buying I/O solutions from vendors like ADI can reduce time to market, save money, and free up valuable resources. However, it is natural to think that one's own engineers and shops can design and assemble a needed function—especially one that is not apparently very challenging—more cheaply and efficiently. Buying I/O functions from other companies is not pursued because in-house engineers and technicians have the necessary expertise (and are enthusiastic about doing the work themselves), because of unawareness of what's currently available, or because of unpleasant past experiences with vendors. It may perhaps come as a surprise to conduct an analysis and find that buying circuit solutions can actually be regarded as a source of added value to the system-level company.

From the *system designer's* perspective, the prospect of make vs. buy means the need to identify trade-offs. There are many considerations in coming to the right decision. Is making the function in-house really cheaper? Does making the function internally help get the end product to market faster? Can we find a good partner? Can we reach our performance goals? Is there risk in sharing trade secrets with potential partners? The answers to these questions are not easily found.

The *vendor's* perspective is one of business decisions. Opportunities are weighed, and commitments are made based on profit margins, potential for future business, and potential for technological benefits. Vendors must efficiently use their resources to remain viable and competitive. They must be able to offer attractive prices and lead times to be successful in the market.

One thing is certain: to be a successful enterprise, the deal has to benefit both parties (a win-win situation). If this is not the case (for example, sufficient information is not available for proper specifications, corners are cut to reduce cost, vendor capabilities are oversold), one or both parties will suffer pain somewhere in the future of the relationship.

GUIDING THE DECISION PROCESS

The following lists the key considerations in the make vs. buy decision process.

- Cost
- Time to Market
- Internal Resource Allocation
- Intellectual Property Protection
- Partnerships
- Weighing Future Benefits

Cost

Determining the cost in a make vs. buy situation is started by having a full understanding of the real internal cost for making the circuit. With this knowledge, the customer can then approach a vendor with a request for quote for the same function. Determination of internal costs should include material cost, labor costs, engineering and support costs, as well as overhead factors. An important area of cost often overlooked is *lost opportunity*, which occurs when scarce resources are used on projects that are not of primary importance.

The cost of buying the solution from a vendor like ADI can be found by simply requesting a comparable quotation of price for the function at the required performance level. The vendor may offer a standard product, or, as needed, semi-custom—or even custom—products.

Vendors can offer many advantages. In the case of Analog Devices, these advantages include the ability to obtain proprietary silicon solutions from product lines throughout the company at prices less than market cost—and access to inexpensive off shore manufacturing. By combining these two attractive features, the cost of industrial I/O functions made by ADI are among the lowest in the industry. By allowing such a vendor to do the entire design, and manufacturing, an OEM company can save material and labor costs, and free up their design resources. This allows the company's talent to concentrate more on their area of added value, such as system, software, and package design.

The acceptance of purchased I/O solutions is well established over years of experience. Analog Devices, for example, has demonstrated great success with the 3B, 5B, 6B, and 7B series of products. These product lines are accepted industry standards, used by many OEMs for industrial I/O.

Time to Market (TTM)

Time to market is another area where purchasing a solution can offer significant benefits. I/O vendors are geared for quickly turning designs, reducing customer TTM. Quick time-to-market is a key to gaining market share. Buying I/O solutions can give a company a competitive edge by simplifying the market introduction process. Often the desired function is a standard product, with a lead-time of days or weeks.

Buying solutions allows for a company to focus on the many details of designing and introducing a new product without wasting internal resources and spending time on the design, implementation, and testing of the I/O functions.

Internal Resource Allocations

The profitable dedication of internal resources to the right tasks is a key challenge for a successful business. Companies insisting on their own I/O solutions often put themselves at a disadvantage compared to companies solving the same problem by buying similar functions from a vendor. The key point for OEMs to understand is the true nature of their company's value-added; in a system-level company, value is not added by designing and assembling industrial I/O devices that could be purchased.

Intellectual Property Protection

The area of industrial control is specialized. Technological expertise, trade secrets, developed by years of participation in the industrial market, so-called *intellectual property*, is closely guarded. Sometimes the fear of losing this information prevents companies from discussing technological options openly with a vendor who is capable of providing solutions at reasonable cost.

The mechanism for assuring protection is a Non-Disclosure Agreement (NDA), which is agreed to by both parties. Reputable vendors adhere to good business ethics. Sharing of proprietary information is important when partnering with another company. The NDA spells out exactly the information that is deemed proprietary. This legally protects the companies involved from the loss of intellectual properties and rights.

With a signed NDA, both companies can reach a level of comfort and trust with each other. From here, solid business relations and partnerships are launched. This environment facilitates free flow of technical ideas and information and ultimately leads to the best and most cost-effective solutions for technical problems.

Partnerships

ADI realizes that having partners in business (i.e., establishing an ongoing relationship between an OEM and a vendor) can prove very beneficial if the arrangement is set up properly. The topics discussed earlier suggest what it takes to locate, communicate, and deal with a good partner. It's worth repeating that, for a relationship to be successful, the deal has to be good for both companies (a win-win situation for both).

Of course, it is important to understand that competitors (of both the OEM and the vendor) are thinking of the same options.

Nevertheless, within the industrial control market, the use of partnerships is very common.

Weighing Future Benefits

Quality products in the Analog Devices Input-Output Subsystems (IOS) product line have a both a proven record of offering cost effective high performance industrial I/O functions/solutions and the prospect of yet more user advantages going into the future. Many companies have realized benefits through developing a channel with ADI. By buying these products from ADI they can obtain a significant competitive advantage

REAL-WORLD EXAMPLE

We will use a tangible example to illustrate the benefits of buying, instead of making, a circuit function. The assumptions used in this example reflect realistic costs for material, services, and labor in today's market. This example is taken from experiences within ADI's IOS (*Input Output Subsystems*) product line.

The circuit function in question is an isolated, gain-of-1, signal-conditioner module, 1-5 V input to 1-5 V output for the front end of a large process control system. This function is commonly used in the process control industry for isolating terminated (250 ohms) 4-20 mA current-loop signals. The volume projection for establishing the cost of this function is 1000 units/year.

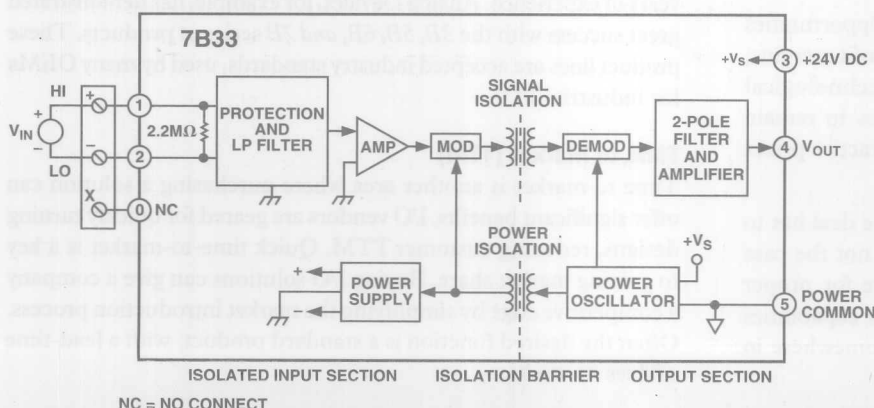
A block diagram of this simple function and a list of key specifications are shown below.

BUY Solution

A recommended option is to purchase the 7B33-01-1 from Analog Devices. The following outlines the consequences of this purchase:

Cost per Unit (100s)	<\$60
Time to Market	4 Week Lead Time
Internal Resource Allocations	Minimal, Principally Component Verification
Intellectual Property Protection	No Risk
Partnerships	ADI, Industry Leader
Future Benefits	Industry Standard Product Format. Other Functions Available

By purchasing the solution, the buyer will spend \$60,000 annually. The product is available in 4 weeks, with lead times reducible to delivery from stock (or just-in-time scheduling), if forecasting is provided. The buyer benefits from ADI's years of experience in applying isolation technologies, cost-effective robust packaging,



KEY SPECIFICATIONS

Accuracy (Calibration Error)	0.10% max
Linearity Error	0.02% max
Common-Mode Rejection	>100 dB
CMV, Continuous	1500 V rms
Input Impedance	>1 megohm
Input Bias Current	100 pA
Bandwidth	100 Hz
Power	+24 V @ <20 mA
Operating Temperature Range	-40°C to +85°C
Input Protection	120 V ac
Agency Certifications	CE, CSA

and standardized families of interchangeable products. Also all ADI products go through extensive reliability, and environmental testing to ensure product quality.

Of particular note is the ability of 7B-Series products to withstand 1500 V rms continuous common mode voltages. This requires special process techniques in fabricating the transformer and the board (PCB)—techniques that have been refined over years of experience with isolation-type products. Finally the product comes with CSA and CE certification.

The user will have to make a small investment of internal resources for the normal acceptance process of verifying that the component (7B33) meets specifications; and of course design verification of overall performance in the system is usual. An additional source of confidence is the knowledge that an industry-standard product from an industry leader is being used in the system, and that other I/O functions are also available in this same product family. Naturally, increased discounts are available if the user's product is successful and quantities increase.

MAKE SOLUTION

The following analysis of the "real" cost to make this function includes an itemization of typical costs for the entire design process, in-house manufacturing, design verification, documentation, material costs, and agency approvals.

In-house engineering and engineering services are estimated using a rate of \$75/hr or \$3000/week, CE and CSA approvals costs are estimated at \$3,000 (although this cost may be subsumed in the costs of agency approval for the overall system). The material cost for this function is about \$25/unit, labor and overhead costs to build the product are \$25/unit.

The cost of the design is estimated as:

Design Engineer	4 Weeks	\$12,000
PCB Layout	3 Weeks	\$ 9,000
Packaging Design	4 Weeks	\$12,000
Documentation	4 Week	\$12,000
Design Verification	8 Weeks	\$24,000
Resource Costs	23 Weeks @ \$3000/Week	\$69,000
Agency Approval		\$ 3,000
Total Development Costs*		\$72,000

*(\$72 per unit for the first 1000 units)

It should be noted that this example assumes a reasonable degree of competence in I/O function design. Achieving an adequate level of performance, price, and reliability may turn out to be difficult for the unaware. Transformer design is a critical area to insure meeting the isolation-voltage specification. For the sake of this demonstration, these key areas are assumed to have been addressed and under control. Further, if this were a signal conditioner that had more design ramifications, such as a linearized thermocouple input module (7B47), the assumed risk becomes even greater.

Once this design is released, it can be manufactured for approximately \$50/unit. This is a \$10/unit "saving" from the purchased solution, or \$10,000 annually. Just to pay back the \$72,000 investment at this rate, the design will have to be in use for about 10 years (\$73,000 is the approximate present value of \$10,000 per year for 10 years at 6%), Considering that this is only break-even, and that the invested capital has not generated any return, this is a poor use of scarce resources by any standard.

A much more desirable and usual expected payback period is 2 years. If only \$20,000 has been paid back on the development costs in 2 years, the \$50,000 that was sunk ($\$50,000/2,000 = \25 per unit) should be added to the manufacturing cost, resulting in a real cost per unit of \$75, or 20% more than the cost of purchasing the devices. Thus, by making the function, the company incurs a large and unnecessary risk, and in fact has lost money because of this unfortunate business decision.

The development schedule for this type of product is approximately 20 weeks. This assumes that every step in the development goes perfectly. Schedules are always an area of risk. It is likely that the design and release of this simple I/O function will be on the critical path for the overall process control system release.

By making instead of buying, this company has put its system release at risk.

The use of internal resources is wasteful and adds an unnecessary opportunity cost. Key resources are wasted on projects that do not add value to a system-level company's products. The evidence of this is the negative return on invested capital. If these resources were applied to *system* design, or some other value-added task that improves efficiency, the company could realize (for example) faster time to market for its end product.

The money put toward designing this simple I/O function is an expense, not an investment.

By choosing to make this product, the systems company has gained some organizational learning, but in a technological and business area that is not on the mainstream of success in their field. In addition an opportunity has been missed to develop more knowledge in their own systems business by spending resources more relevantly.

Making this function is a losing proposition.

Because the company did not look for or choose to "partner", they are at risk from their competition. The company will realize no benefit in the future because of this missed opportunity to partner with a leading company in a joint venture.

Here is the bottom line of the "Make" solution results:

Cost	\$75
Time to Market	20 weeks, on critical path
Internal Resource Allocations	Significant and high risk
Intellectual Property Protection	N/A
Partnerships	None
Future Benefits	Little, if any

SUMMARY

The above example is a realistic assessment of a typical choice situation in today's market place. It shows that there is a trap in deciding to "Make" circuit solutions instead of buying the solution from a substantial, experienced, and well-qualified vendor with a proven product line. In all the key factors that contribute to making the decision, "Make" is shown to be riskier and more costly. The lesson to be learned is that designers and decision makers in enterprises must at all times be aware of what their real business is and carefully think through the consequences of make-or-buy decisions that "re-invent the wheel." Decisions that fail to take advantage of existing products, and the customization capabilities of their manufacturers, may well weaken the company in the long term. The discipline of a bias towards "Buy" for high-performance circuit solutions can be a valuable asset. ▀

New Fellows

Three new Fellows were named in 1999: Bob Adams, Frank Murden, and Jake Steigerwald. Fellow, at Analog Devices, represents the highest level of achievement that a technical contributor can achieve, on a par with Vice President. The criteria for promotion to Fellow are very demanding. Fellows will have earned universal respect and recognition from the technical community for unusual talent and identifiable innovation at the state of the art. Their creative technical contributions in product or process technology will have led to commercial success with a major impact on the company's net revenues.


Attributes include roles as mentor, consultant, entrepreneur, organizational bridge, teacher, and ambassador. Fellows must also be effective leaders and members of teams and in perceiving customer needs. This trio's technical abilities, accomplishments, and personal qualities well qualify them to join Woody Beckford (1997), Derek Bowers (1991), Paul Brokaw (1979), Lew Counts (1983), Barrie Gilbert (1979), Roy Gosser (1998), Bill Hunt (1998), Jody Lapham (1988), Chris Mangelsdorf (1998), Fred Mapplebeck (1989), Jack Memishian (1980), Doug Mercer (1995), Mohammad Nasser (1993), Wyn Palmer (1991), Carl Roberts (1992), Paul Ruggerio (1994), Brad Scharf (1993), Mike Timko (1982), Bob Tsang (1988), Mike Tuthill (1988), Jim Wilson (1993), and Scott Wurcer (1996) as Fellows.

BOB ADAMS

Bob Adams joined Analog Devices in 1989 as Manager of Analog Technology. He quickly established himself as a leader within the ADI technical community, particularly in the area of high-resolution sigma-delta conversion. Since then he has led the design of numerous digital-audio products, including sigma-delta A/D and D/A converters, sample-rate converters, and codecs. Bob developed the first multi-bit sigma-delta converters in the industry, using two new patented algorithms. He also developed the industry's first monolithic asynchronous sample-rate converters, the AD1890 family. As Design Manager for the Digital Audio Group, he is driving developments in conversion and surround codecs that are establishing Analog Devices as a leader in the consumer audio industry.

Bob was graduated from Tufts University in 1976, and before joining ADI he was Director of Audio Research at dbx Inc. Bob is a Fellow of the AES and holds more than 15 patents, mostly relating to audio signal processing. He recently received the "Best Paper" award at ISSCC for his 1998 presentation of a new DAC architecture.

In his spare time, Bob plays the saxophone in various bands (including the "Idle Tones," an infamous ADI band). At home, he and his wife, Susan, are kept very busy with their two children.



FRANK MURDEN

Frank Murden joined ADI's Greensboro operation in 1983, after he had earned an MSEE from the University of Arizona. In his career at Analog Devices, he has developed a great many high-performance state-of-the-art A/D converters, starting with hybrid circuits and moving to IC types as the technologies improved—improvements that were helped in no small measure by his valuable feedback to process-development teams.

Examples of his designs include the early (and long-unchallenged) AD9014, 14-bit, 10 MSPS ADC, followed by a series of converters with speeds and resolutions leapfrogging one another, from the 12-bit, 30 MSPS AD9026 family to the 14-bit, 80 MSPS AD6644, for use in multicarrier, multimode receivers ("software radios").

Frank has published technical articles, presented at ISSCC, co-authored MTT papers, and received a number of patents on receiver architectures, ADC architectures, and circuit topologies. He is also a frequent peer reviewer of papers for IEEE. His recent work on diversity-receiver chipsets, such as the AD6600, has resulted in their incorporation in designs by many of the world's major communications equipment manufacturers.

Apart from his design efforts, Frank has served as mentor to many young engineers and has developed close relationships with engineers in widely separated entities (both organizationally and geographically), within Analog Devices and among our customers.




JAKE STEIGERWALD

Since his arrival at Analog Devices in 1992, Jake Steigerwald has been a key player in the development of a number of high-speed complementary bipolar and RF-oriented integrated-circuit fabrication processes, and in the improvement of existing processes. Areas of particular interest for his expertise have been yield improvement, low-gate-current JFETs, exploration of higher-voltage high-speed processes, and exploratory work in the early stages of silicon-germanium process development.

Prior to joining ADI, Jake worked at National Semiconductor and, more recently, Commodore Semiconductor. He has a BSEE from Clarkson and has taken graduate courses at RPI, Yale, Villanova, University of Utah, and Boston University.

Jake's forte is process and device *simulation*, which he has used to great advantage in accumulating understanding as a basis for process innovation. These modelling techniques have resulted in big savings in experimental silicon and in rapid activation of new and improved semiconductor fabrication processes. He has been a mentor to ADI process engineers at all levels, a sought-after teacher of device physics, and a consultant on such topics as regional boundary approaches to heterojunction simulation.



MORE AUTHORS [Continued from page 2]


Jim Doscher (page 27) manages the consumer/industrial business unit at the Micromachined Product division of Analog Devices and concentrates on the creation of new markets for ADI's micromachining technology. Before that, he was a Marketing Manager for the Sensor Products group. Before that, Jim worked on quality improvement programs and facilitated the development of a corporate-wide total-quality management program at ADI. Earlier he had worked as a Product/Test Development Engineer for new products and was a product marketing engineer with the Linear Products group. He joined Analog Devices Semiconductor after graduating from Dartmouth College with an AB and BE.



Bill Englemann (page 58), as the Northeast Region Distribution Manager for our North American sales team, is responsible for all sales through electronic distributors in New England, Upstate New York, and the Metro New York City area. Bill has a BSEE from Worcester Polytechnic Institute and has held applications, marketing, and sales management positions with several divisions of Analog Devices. Outside of work he serves as a corporate agent for the WPI alumni organization and enjoys carpentry and making home improvements.



Dan Fague (page 3) is the RF systems Applications Manager in ADI's Communications division (Wilmington, MA). Since joining the company in 1995, he has been responsible for both RF system architecture design and reference radio design for ADI's RF chip sets. He holds a BSEE from Gonzaga University and an MSEE from the University of California at Davis. When not working, Dan can be found at the beach or playing volleyball (or both).



Rich Ghiorse (page 67) is a Senior Product Engineer in ADI's Motion Control IC group. In his 19 years at ADI, he has designed products and test equipment, working with isolation amplifiers and signal conditioners (3B, 5B, 6B, and 7B Series). He holds BSEE and MSEE degrees from Northeastern University. An avid golfer, he is happily married, with 3 children.



Michal Gwozdz (page 54) has worked at P.E.P. ALFINE (ADI's Representative in Poland) since 1997 as a DSP Field Application Engineer. Starting his career in telecommunication engineering, he became a scientific worker, specializing in computer control

circuits for power electronic devices for 6 years at the Technical University of Poznan (Institute of Electrical Engineering). He received a Ph. D. from the same University, with a major interest in power-electronics filters. At ALFINE, he continues to work on highly advanced control algorithms of power electronics equipment using DSP. His publications include more than thirty scientific papers, articles and patents in his specialty. He is married, with two sons, 10 and 17. His greatest passions are mountain tourism and travelling.



John Markow (page 62), in 1999 was an Applications Engineering Intern with ADI's General Purpose Converters group in Wilmington, MA. Currently, he is working towards his MSEE in Integrated Circuits at the University of California at Berkeley. He has a BSEE from Worcester Polytechnic Institute and completed his senior project, a dspConverter™ evaluation kit, while working at ADI, Limerick, Ireland, in 1998. In his spare time, John enjoys playing trombone, hiking, travel, and tennis.



Mike McKeon (page 67), a 29-year veteran at ADI, is a Technical Operation Manager in the Motion Control IC group, responsible for Industrial I/O and IC Isolation Amplifier product lines. He also has foundry management responsibilities for our offshore manufacturing site in the Philippines. He has a BS in Industrial Management from Boston State College (now U. Mass. Boston). In his spare time, he enjoys swimming, golf, and pool.




Eamon Nash (page 30) is an Applications Engineer, now at ADI's Advanced Linear Products group, and formerly of the Central Applications group. He was graduated from the University of Limerick, Ireland, with a Bachelor of Engineering degree in Electronics. His spare-time activities include tennis, squash, and rollerblading.




Paul O'Brien (pages 9-22) is a Senior Test Engineer in the RF group at Analog Devices, Limerick. He has BEng and MEng degrees in Electronic Engineering from the University of Limerick. Before coming to Analog Devices, Paul had worked on GSM transceivers, interactive cable TV, and multiple-access radio systems at Lucent Technologies. At play, he enjoys fishing and golf.



Bill Odom (page 41) is an Applications Engineer for ADI's High-Speed Converter group in Wilmington, MA. He has a BSET from the New York Institute of Technology, at Old Westbury, NY. During the last 20 years Bill has been involved in the design and application of A/D converters, at ILC Data Devices and Signal-Processing Technologies, Inc. He enjoys camping and entertaining his family.



Albert O'Grady (page 60) is a Senior Engineer with the Applications group in Limerick, Ireland, where he provides customer support for high-resolution, low-bandwidth sigma-delta A/D converters and general-purpose D/A converters. He is also involved in new-product definition for these products. Albert holds a BEng from the University of Limerick. In his spare time, he enjoys reading and plays badminton and tennis.



Rasekh Rifaat (page 6) is a Systems Application Engineer for the Wireless Infrastructure business group within the DSP division. He specializes in algorithm implementation for 3G applications using ADI's TigerSHARC Static Superscalar DSP. Rasekh holds a Master's degree in Computer Engineering from the University of Manitoba, Winnipeg, Canada. In his spare time, Rasekh enjoys movies, soccer, biking and traveling. He has traveled extensively in Canada and Germany.



Matt Smith (pages 44-53) is a staff Applications Engineer at the Limerick, Ireland, facility, with responsibility for interface and hardware-monitoring products. He holds a BEng from the University of Limerick. His leisure interests include playing squash, motor maintenance, and woodworking.



Harvey Weinberg (pages 23-26), is the Applications Engineer for the Analog Devices Micromachined Products division, principally for non-automotive markets, for the past two years. He has a BEng from Concordia University (Montreal, Canada), and he has 12 years of design experience in precision analog and microcontroller circuits. Before joining ADI, he was an Application Engineer for Future Electronics, an electronics distributor, where he provided applications engineering support for several major product lines, including Analog Devices.



Worth Reading

AN AUTHORITATIVE, PRACTICAL, READABLE NEW DSP BOOK Available FREE on line from Analog Devices¹

Digital signal processing is one of the most powerful technologies that will shape science and engineering in the twenty-first century. The problem is, if you can't understand it, you can't use it! A new book available from ADI can help: *The Scientist and Engineer's Guide to Digital Signal Processing*, by Steven W. Smith, Ph.D. The best part is that it's free; you can download individual chapters and thus the entire book in electronic form (PDF files). A print version in soft cover is available for purchase from Analog Devices.

Unlike most books on DSP, which contain extensive mathematics (a requirement for use in Electrical Engineering courses), this book is for scientists and engineers who need a practical understanding of DSP, but who don't have the time or background to learn the rigorous theory. For those new to the field, its informal style and clear explanations make it an ideal introduction; and it's an excellent reference and refresher for the more experienced. The reader is assumed to have no previous knowledge of DSP algorithms or hardware, but it helps to have a basic understanding of analog and digital electronics. Unlike most of the DSP literature that uses extensive complex number theory, this book explains all of the basic techniques using only simple algebra. Complex numbers are treated as an advanced topic, something to be learned after the fundamentals are understood. The software programs in the book are written in a very elementary form of BASIC, making them easily understood by anyone with even minimal programming experience.

The book covers a wide range of topics, beginning with elementary techniques and progressing to quite sophisticated algorithms. The first section, entitled Foundations, deals with statistics and noise, how A/D and D/A conversion affect information content, and software concerns—such as number representation and round-off error. This is followed by a section on Fundamentals, describing the classic techniques of convolution and Fourier analysis. The third section, Digital Filters, is expected to be the most widely used portion of the book. The power and flexibility of digital filters is one of the key reasons that DSP has become so popular. This section describes the characteristics and design of a wide variety of digital filters, including FIR and IIR, as well as FFT-convolution techniques. Most important, the design methods presented are straightforward, almost cookbook, without the excessive mathematics usually associated with filter design.

The fourth section of the book, Applications, is a diverse collection of specialized techniques, including: data compression, neural networks, audio processing, and image processing. Two chapters in this section are devoted to DSP hardware, featuring the Analog Devices SHARC[®] family of digital signal processors. The fifth and last section of the book covers Complex Techniques, providing an introduction to the mathematics that theoretical DSP is based on.

The book can be downloaded from our DSP site (*free*) at http://www.analog.com/industry/dsp/dsp_book/. To purchase the book in soft cover, in the U.S.A. call 1-800-262-5643 and select option 3; in Europe call +49-89-76903-0. The soft-cover international standard book number (ISBN) is 0-9660176-4-1. Visit the author's Web site at: <http://www.dspguide.com/> for hardcover sales, information about the author, and links to other DSP sites. ■